NBER WORKING PAPER SERIES

IN THE FED'S MIND

Ali Kakhbod
Amir Kermani
Bernardo Maciel

Working Paper 35016
http://www.nber.org/papers/w35016

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2026

We thank Michael Bauer, Philip Bond, Anna Cieslak, Stephen Hansen, Jonathan Hazell, Lukas Kremens, Emi Nakamura, Ricardo Reis, Jón Steinsson, and seminar participants at the Berkeley Macro Lunch Seminar, LSE Money-Macro Seminar, University of Washington, and UCL for very helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

In the Fed’s Mind
Ali Kakhbod, Amir Kermani, and Bernardo Maciel
NBER Working Paper No. 35016
March 2026
JEL No. E31, E43, E52, E58

ABSTRACT

Does the Federal Reserve react to all inflation equally? We systematically analyze FOMC meeting records from 1937 to 2025 to construct meeting-level measures of the Fed’s real-time attribution of inflation to demand and supply pressures. We document substantial variation in these narratives over time and show that, since Volcker, the Fed has responded more aggressively to perceived demand-driven inflation. Consistent with this asymmetry, supply pressures have more persistent effects on realized inflation, while demand pressures' impact dissipates quickly. Financial markets also reflect this distinction: demand imbalances primarily move risk-neutral yields, whereas supply imbalances raise term premia and equity risk premia.

Ali Kakhbod
University of California, Berkeley
Haas School of Business
akakhbod@berkeley.edu

Bernardo Maciel
University of California, Berkeley
Haas School of Business
bernardo_maciel@berkeley.edu

Amir Kermani
University of California, Berkeley
Haas School of Business
and NBER
kermani@berkeley.edu

Introduction

The sharp rise in inflation following the Covid-19 pandemic reignited a long-standing debate in macroeconomics: what are the underlying causes of inflation, and how should monetary policy respond? In particular, the distinction between supply-driven and demand-driven inflation has taken center stage. In standard macroeconomic models, supply and demand shocks have markedly different implications for optimal monetary policy. A central bank that reacts aggressively to supply-driven inflation risks deepening a slowdown, while failing to respond to demand-driven inflation risks entrenching price pressures. Understanding how policymakers interpret the sources of inflation and how they react to them in real time is thus essential for evaluating and predicting monetary policy decisions.

Empirically, however, distinguishing between supply-driven (including cost-push shocks) and demand-driven inflation in real time is notoriously difficult. Traditional approaches usually rely on structural macroeconomic models. While valuable, these techniques typically require strong identifying assumptions and often rely on data that were not available to policymakers at the time of decision-making. Moreover, they reveal how inflation can be decomposed from an econometrician's perspective—not necessarily how central banks themselves viewed the shocks as they unfolded.

We build on the narrative approach pioneered by Friedman and Schwartz (1963) and expanded by Romer and Romer (1989, 2023) to systematically analyze FOMC meeting records from 1937 to 2025. Using a two-stage extraction and classification procedure grounded in macroeconomic sign restrictions, we construct a high-frequency panel of the Federal Reserve's real-time attribution of inflation to demand and supply components. We first document substantial time variation in the Fed's assessment of the relative importance of demand and supply forces in explaining inflation, with supply-related factors playing a particularly prominent role in the late 1970s and during the post-Covid inflation surge. We then show a pronounced policy asymmetry: interest rates respond more aggressively to perceived demand-driven inflation than to supply pressures, especially since Volcker. Consistent with the Fed's policy response, we find that demand-driven inflation pressures fade quickly, while supply-side pressures persist in realized inflation. Finally, we show that the Fed's diagnosis is reflected in financial markets, with demand imbalances primarily raising risk-neutral yields, while supply pressures predominantly increase bond and equity premia and decrease equity prices.

Recent advances in large language models (LLMs) allow us to extract and classify economic content from the full corpus of FOMC meeting records. These records document the discussions underlying monetary policy decisions and provide a direct window into policymakers' real-time interpretations of macroeconomic conditions. Our approach proceeds in two steps. First, we use LLMs to identify statements that explicitly cite reasons for inflation or disinflation. Second, we classify each reason as demand-driven or supply-driven using a theory-based criterion: whether the factor is expected to move inflation and output in the same or opposite directions. We then aggregate these classifications into meeting-level imbalance measures that capture the normalized net share of inflationary versus disinflationary remarks within each category.

These measures closely track well-known historical inflation episodes, such as supply-driven pressures during the 1970s oil shocks and weak demand during the Great Recession. After Covid, the measures capture a sharp transition from disinflationary risks in early 2020 to inflationary supply factors in 2021–2022, driven by global supply-chain disruptions and energy price increases following Russia's invasion of Ukraine.

To frame the empirical analysis, we next develop a simple New Keynesian model with demand and supply shocks that maps the Federal Reserve's perceived diagnosis of inflation into monetary policy, realized inflation, and asset prices. Under discretionary optimal policy, the Fed fully offsets demand disturbances but partially accommodates supply disturbances because of the inflation–output trade-off. This asymmetry carries directly into asset prices. Demand shocks primarily affect nominal yields through the risk-neutral channel, reflecting changes in the expected path of policy rates. Supply shocks, by contrast, reduce output and consumption, raising the marginal utility of consumption in adverse states and inducing a positive covariance between nominal bond payoffs and the stochastic discount factor. As a result, supply shocks raise primarily term premia rather than expected short rates, while also increasing equity risk premia. The model delivers three clear and testable predictions: the Federal Reserve responds more aggressively to perceived demand shocks than to supply shocks; realized inflation is more strongly driven by supply disturbances; and asset prices reflect this asymmetry, with demand shocks moving risk-neutral yields and supply shocks raising term premia and equity risk premia.

Motivated by the model's predictions, we estimate an augmented Taylor rule and find a clear asymmetry in monetary policy responses: the Federal Reserve reacts more aggressively to perceived demand-driven inflation while partially looking through supply pressures. This differentiation across inflation sources emerges during the Volcker era and remains a defining feature of U.S. monetary policy in the decades that follow. Consistent with this asymmetric response, we find that demand-driven inflation pressures dissipate relatively quickly, whereas supply pressures have more persistent effects on realized inflation. As a result, observed inflation outcomes are driven predominantly by supply-side disturbances.

We next take the model's asset-pricing predictions to the data. Using Adrian, Crump, and Moench (2013) yield decompositions around FOMC announcements, we find that demand imbalances primarily raise the risk-neutral component of yields, consistent with a tighter expected policy path, whereas supply imbalances primarily increase term premia. In high-frequency windows, supply imbalances are associated with lower equity prices and higher volatility, while demand imbalances elicit a much weaker equity response. These patterns align with the model's asset-pricing implications: supply shocks cause nominal bonds to comove with equities and therefore command a premium, whereas demand shocks do not. Standard pre-meeting macroeconomic indicators explain only 30 percent of the variation in the Fed's diagnosis narratives, indicating that FOMC discussions contain information beyond publicly available data. Controlling for macro news maintains the link between our imbalance measures and movements in risk-neutral rates and term premia.

Contribution to the Literature. Our paper connects four literatures: inflation decomposition, policy response estimation, text-based monetary policy analysis, and monetary policy shocks.

First, we contribute to the empirical literature that decomposes inflation into demand and supply components. Existing approaches fall into two groups. The first approach estimates structural macroeconomic models and uses historical shock decompositions to attribute inflation movements to aggregate demand, productivity, and cost-push shocks (e.g., Smets and Wouters (2007)). More recently, a growing literature exploits input–

output matrices and production networks to disentangle aggregate demand forces from sectoral supply constraints and bottlenecks, quantifying how shocks propagate across industries and contribute to inflation dynamics (e.g., Baqaee and Farhi (2022), di Giovanni et al. (2023)).

More closely related to our approach, the second group decomposes inflation into demand and supply drivers using sign restrictions. Eickmeier and Hofmann (2022) propose a quarterly structural factor model identified with sign restrictions and estimated on a large set of inflation and real activity indicators, while Shapiro (2024) also relies on sign restrictions but implements them in a disaggregated setting, using sectoral price and quantity information from the monthly Personal Consumption Expenditures index to construct demand- and supply-driven inflation measures. Unlike these approaches, ours requires neither a structural model nor data unavailable to FOMC participants at the time of the meetings.

Second, we contribute to the literature on the estimation of the monetary policy rule. Taylor (1993) showed that a simple linear rule not only fits the Fed's decisions during the late 1980s and early 1990s but also serves as a normative benchmark for anchoring long-run inflation expectations. Subsequent work by Clarida, Galí, and Gertler (2000), Cogley and Sargent (2001, 2005), and Kim and Nelson (2006) document significant time variation in the Fed's policy function, with a markedly stronger response of monetary policy to inflation since the Volcker era. Orphanides (2002, 2003) instead emphasize that real-time mismeasurement of macroeconomic variables, particularly the output gap, played a central role in the Great Inflation of the 1970s, rather than changes in the policy rule itself. More recently, Bauer, Pflueger, and Sunderam (2024) focus on time variation in perceptions of monetary policy and show that professional forecasters' expectations also imply substantial changes over time in how monetary policy is perceived. Stein and Sunderam (2018) provide a complementary channel: when the Fed holds private information about its target rate, aversion to bond-market volatility induces gradualism, and each policy move reveals information to markets, amplifying the response of long rates.

We complement this literature by showing that changes in how the Fed diagnoses the sources of inflation, not only time variation in the reaction function, generate differential policy responses to similar deviations of inflation from target. Our augmented Taylor rule results are most closely related to Hofmann, Manea, and Mojon (2024), who build on the demand and supply decomposition of Shapiro (2024). Our finding that the Fed responds differently to inflation depending on its underlying source is consistent with a key prediction of New Keynesian models, namely that optimal monetary policy may partially look through inflation driven by supply shocks (e.g., Clarida, Galí, and Gertler (2000), Guerrieri et al. (2021), and Beaudry, Carter, and Lahiri (2023)), particularly when policymakers are less concerned about their reputation (Nakamura, Riblier, and Steinsson, 2025).

Third, we contribute to the growing literature that uses natural language processing and large language models to study monetary policy. Lucca and Trebbi (2009) compute "hawkishness" scores from FOMC statements, while Gorodnichenko, Pham, and Talavera (2023) and Curti and Kazinnik (2023) construct measures based on the tone and facial expressions of the Fed Chair during post-FOMC press conferences. Hansen and Kazinnik (2023) use ChatGPT to parse and classify the policy stance of FOMC announcements and compare the results with human classifications. Using FOMC transcripts, Hansen, McMahon, and Prat (2017) study the impact of transparency on internal FOMC discussions; Cieslak et al. (2023) measure inflation uncertainty and show that higher uncertainty is associated with a tighter policy stance; Cieslak and McMahon (2023) show that more hawkish stances are associated with lower risk premia; and Shapiro and Wilson (2022) infer time-varying central bank objectives from policymakers' statements. Aruoba and Drechsel (2024) show that the text of Federal Reserve staff documents contains information beyond numerical forecasts. We contribute to this literature by showing that the Fed's internal deliberations contain information about policymakers' real-time assessments of the sources of inflation, with important implications for monetary policy and asset prices.

Fourth, we contribute to the literature on monetary policy shocks. A key challenge in this literature is that different shock constructions are often associated with distinct implications for asset prices and aggregate outcomes (Brennan et al., 2024). Our results show that, consistent with the predictions of Cieslak and Pflueger (2023), asset price reactions to monetary policy announcements depend critically on the Federal Reserve's assessment of the underlying sources of inflationary pressures. In particular, supply-driven inflation shocks have a large impact on term premia, whereas demand-driven shocks primarily affect risk-neutral rates. Accordingly, variation in empirical findings across different monetary policy measures may partly reflect differences in how these measures load on demand- and supply-driven inflation pressures.

The remainder of the paper is organized as follows. Section 1 discusses the data. Section 2 describes our methodology and explains how we use a large language model to identify and count the Fed's perceived causes of inflation over time. Section 3 presents a simple model and the intuition for the results. Section 4 presents the empirical results. Section 5 concludes.

1 Data

We combine FOMC meeting texts with standard macroeconomic series; Appendix Table D.1 reports sample periods and observation counts. Our baseline corpus covers 766 FOMC meetings from 1937 to 2025, drawing on minutes and transcripts.[1] We restrict estimation to the post-Accord period (1960 onward), where PCE inflation data are available and the policy framework is coherent.[2]

We assign each meeting the macro value observed in its calendar month, so all variables are at the FOMC-meeting frequency. Core and headline PCE inflation and the effective federal funds rate are from FRED. During zero-lower-bound episodes, we replace the funds rate with the shadow rate of Wu and Xia (2016), which extends the Black (1995) option-based framework to a term-structure model.[3] The Adrian, Crump, and Moench (2013) term premium and risk-neutral yield decomposition is from the New York Fed, and the Bauer and Swanson (2023*b*) monetary policy surprises and macro news data come from the San Francisco Fed. Additional data sources used in specific analyses are described in the relevant sections.

[1]The format of the minutes changes over time. Before 1967, we use the "Historical Minutes" from the Fed website; the "Memoranda of Discussion" between June 1967 and March 1976; the transcripts between March 1976 and December 1992; and the modern "Minutes" from 1993 onwards. We do so because minutes are more informative about economic conditions after 1993 and less so before that. Using minutes for the most recent period also allows us to avoid the five-year embargo for the transcripts.

[2]See Eichengreen and Garber (1991) for a detailed discussion of monetary policy under the pre-Accord regime, which operated as an implicit price-level target and, therefore, required pegged nominal yields.

[3]We define the zero-lower-bound period as the two spans when the Federal Reserve held the target federal funds rate in the 0–0.25 percent range: first from December 16, 2008 to December 15, 2015, and again from March 15, 2020 to March 16, 2022. We obtain the shadow rate series from the Atlanta Fed's website.

2 Methodology

2.1 Inflation attribution procedure

We build a measure of the Fed's own diagnosis of inflation by coding every FOMC meeting record from 1937 to June 2025. A large language model (LLM) identifies all reasons that participants cite as causes of inflation in each set of minutes. Unlike proxies based on ex-post price indices, our measure reflects what policymakers said when they set policy. The procedure has two stages.

First Stage. The first stage identifies the reasons the Fed cited as causes of inflation at the time of each meeting. We feed each set of minutes into OpenAI's GPT-5 and ask it to list all reasons that participants explicitly cite as structural causes or underlying factors driving inflation. A structured-object API programmatically populates each response field, turning thousands of pages of minutes into a clean panel that researchers can merge, rerun, and extend.

To ensure robustness and the quality of the results, we explicitly instruct the LLM not to include self-referential answers such as "increasing inflation" or "inflation concerns" and to anchor every classification in a direct quotation from the text. For each identified reason the model records: (i) a concise label ("reason"), (ii) a supporting quotation ("explanation"), and (iii) whether participants judged the reason to push prices up ("I"), down ("D"), or gave no clear direction ("na"). It also records the speaker and assigns an importance score. The full prompt and additional details are described in Internet Appendix IA.7.1.

Second Stage. In the second stage, we feed each reason–explanation pair back into the LLM and classify it as driven by demand, supply, expectations, monetary policy, or indeterminate/unclear.

We provide the model with an objective benchmark: the theoretical sign restrictions implied by demand and supply shifts. A rightward shift of aggregate demand raises both prices and output; a leftward shift of aggregate supply raises prices but lowers output.[4]

[4]One concern is that including the reason label, in addition to the explanation passage, may bias the model toward using its knowledge of historical events rather than the text itself. For example, labels such

Two additional categories complete the taxonomy. Expectations captures statements about shifts in inflation expectations; monetary policy isolates references to the policy stance as the channel influencing inflation.[5] Ambiguous or tautological pressures (e.g., "prices increased" without reference to a structural cause) are classified as "other." The prompt is fine-tuned to address edge cases prone to misclassification. The complete prompt and additional details are provided in Internet Appendix IA.7.2.

Separating extraction from classification serves a clear purpose: the first step is a straightforward natural language processing task, while the second asks the model to apply textbook reasoning linking the shock to its effects on output and inflation. Because the classifier sees only the local reason–context pair, not the full minutes, this separation limits both hallucination and look-ahead bias.[6] The result is a panel of demand and supply reasons, each tagged as inflationary or disinflationary, that tracks the Fed's stated view of inflation's causes over almost nine decades. Internet Appendix Figure IA.1 summarizes the full procedure.

Illustrative examples. To illustrate, we consider the discussion from the FOMC meeting on September 16, 2008 – the day after Lehman's collapse. Contractionary reasons dominate: of the twelve reasons the LLM identifies, nine point to falling inflation.

Consider the reason labeled "weakening labor market," with the records noting that "Participants 'generally thought that the outlook for inflation had improved, mainly re-

as "Covid pandemic disrupting supply chains" or "Energy prices going up as a result of the oil shock" may prompt the model to rely on prior knowledge of these episodes. For robustness, we repeat the second-stage classification using only the explanation, omitting the reason label. The resulting measures are highly correlated with those from our main specification. We nonetheless retain the baseline approach because a single passage often contains multiple reasons, and the label helps the LLM identify which specific reason we refer to in that instance. We discuss additional robustness tests against look-ahead bias in Internet Appendix IA.3.

[5] We further instruct the LLM to exclude generic references to anchored inflation.

[6] A general concern when using LLMs for measurement is look-ahead bias: because the model is trained on text written after most of the meetings in our sample, it may implicitly use ex-post information when interpreting historical passages. Our baseline design already mitigates this risk by separating (i) a first-stage extraction task from (ii) a second-stage classification task, and by feeding the classifier only the local reason and context passage rather than the full minutes. As an additional robustness check, we implement an "entity neutering" procedure following Engelberg et al. (2025): we iteratively mask and paraphrase each reason–context pair until an adversarial LLM can no longer infer the meeting date (within a two-week window), and then rerun the second-stage classification and downstream analysis on the neutered text. The resulting demand/supply series are highly correlated with the baseline measures, and our main regression conclusions are unchanged. See Internet Appendix IA.3 for more details.

flecting ... the weakening of the labor market." A weakening labor market shifts the IS curve left, lowering output; via the Phillips curve, lower slack also reduces inflation pressure. Because output and inflation move in the same direction, the model classifies this as a contractionary demand pressure.

The model records a brief, text-anchored rationale (the "scratch" field) explaining the sign assignment (e.g., "weaker consumption → output ↓ and inflation ↓"). The rationale serves two purposes: it improves consistency by enforcing theory-based sign rules and gives researchers a transparent trail to verify decisions, especially in edge cases.[7] Full examples of the model's output are provided in Internet Appendix IA.6.

The next excerpt of the LLM output is from the June 15, 2022 meeting, amid lingering supply-chain disruptions from China's lockdowns and the Russian invasion of Ukraine. The highlighted reason, "ongoing supply bottlenecks," refers to supply-chain disruptions and rising energy and commodity prices. The scratch field traces the causal chain: bottlenecks constrain output and raise costs, pushing inflation up and output down. Because inflation and output move in opposite directions, the LLM classifies this reason as supply-driven. The full output is also available in Internet Appendix IA.6.

2.2 Summary Statistics

Appendix Table D.2 lists the most common reasons by category. Common demand reasons link the output gap to inflation, such as tight labor markets and economic slack. Common supply reasons cite cost-push factors (e.g. energy prices) and productivity. Expectations reasons describe both rises and falls in inflation expectations.

It is worth noting that reasons related to wage increases appear under both demand and supply. This is because our prompts direct the LLM to distinguish between demand- and supply-related wage increases. As a result, strike-induced or statutory wage hikes are classified as supply pressures, while broad labor-market strength is tagged as demand. This disambiguation is possible because, in the second stage, the model receives both the label and the accompanying passage in which that reason was discussed; thus, identical labels can be mapped to different classifications when the surrounding context differs.

[7] Human auditing of model outputs is a central recommendation in the text-analysis literature (Gentzkow, Kelly, and Taddy, 2019). Moreover, eliciting explicit chain-of-thought reasoning has been shown to substantially improve the quality of large language model inference (Wei et al., 2022).

Concretely, explanations that point to factors such as contracts, strikes, and minimum-wage changes produce a supply classification, while references to tight labor demand or hiring pressures lead to a demand classification.

The Committee's views within a meeting often offset one another: one member stresses a tight labor market pushing prices up, while another might cite weak consumer spending pulling them down. To better understand the distribution of the Fed's overall assessment, we count inflationary and disinflationary remarks separately. Let $N_{x,t}^{(+)}$ and $N_{x,t}^{(-)}$ denote, respectively, the number of reasons on date t from category x that the text describes as raising or lowering inflation, where x can be demand, supply, or expectations. These raw counts summarize the Committee's perceived balance of inflationary and disinflationary perspectives for each category and in each meeting. Table 1 reports summary statistics for all six series and for their totals.

Table 1: Summary Statistics for Reason Effects

Component	Direction	N	Mean	SD	Min	P25	Median	P75	Max
Demand	−	766	0.730	1.077	0	0	0	1	6
	+	766	1.287	1.528	0	0	1	2	9
Supply	−	766	1.484	1.717	0	0	1	2	9
	+	766	2.962	2.653	0	1	2	5	20
Expectations	−	766	0.351	0.688	0	0	0	1	4
	+	766	0.477	0.836	0	0	0	1	5
Total		766	7.291	4.427	0	4	7	10	37

Notes: For each FOMC meeting, we count the number of demand, supply, and expectations reasons described as inflationary (+) or disinflationary (−). The table reports cross-meeting summary statistics for each series and for their sum (Total). All entries are counts of reasons.

Across meetings, the median lists seven tagged reasons. Two facts emerge. First, inflationary tags outnumber disinflationary tags. Second, supply reasons occur roughly twice as often as demand reasons. This pattern reflects our classification rule: we tag a reason only when the passage explicitly links it to inflation. Discussions about demand drivers, such as employment or overall conditions, are excluded unless the text connects them to inflation, even though standard models imply such links indirectly. This discipline depresses demand counts more than supply counts. Including output discussions as implicit inflation reasons leaves the main conclusions unchanged.

Internet Appendix Figure IA.2 plots the total number of reasons we extract for each meeting. The dots mark individual meetings; the black curve is an exponential moving average to capture longer trends. Between the 1960s and the 1990s, the number of reasons per meeting remained remarkably stable, although with a large variance. After the 1990s, this number has increased slightly over time. This is despite the fact that the modern minutes, which are the source used from 1993 to 2025, are much shorter than transcripts, memoranda of discussion, and most of the historical minutes used in the earlier sample.

Internet Appendix Figure IA.3 examines composition by filling a cell whenever a demand, supply, or expectations reason is mentioned, yielding a time-series indicator of each category's presence. Inflationary demand reasons occur throughout the sample but become scarce around the Great Recession; disinflationary demand reasons concentrate in recessions. Inflationary supply reasons also appear across the full period, and more frequently during the 1970s oil shock, the energy-related shocks preceding the Great Recession, and the post-Covid period in 2021–2022. Disinflationary supply reasons arise mainly in the 2010s as energy prices fell, especially after 2015 during the oil glut.

2.3 Post-Accord Narratives

Figure 1 shows net annual totals of demand, supply, and expectations remarks (left axis), with the federal funds rate and inflation on the right axis.[8] We walk through the key episodes below.

Eisenhower recession (1957 – 58). In 1957, the Fed tightened to cool earlier price pressures as autos, steel, and housing slumped, compounded by an overseas liquidity squeeze and the influenza pandemic. Minutes record disinflationary forces—"weak sales", "rising joblessness", "inventory liquidation"—alongside inflationary ones—"defense spending", "easier credit", and restocking after the policy pivot. The Fed's quick rate cuts and fiscal stimulus ended the slide by mid-1958, but price pressures resurfaced in 1959 (Romer and Romer, 2002).

[8] Appendix Figure C.1 repeats the same plot, but separately for reasons increasing and decreasing inflation. Internet Appendix Figures IA.4, IA.5, and IA.6 present snapshots of the same figure on three historical periods: the 1960s, the 1970s, and the Covid pandemic.

Figure 1: Net Aggregated Reason Counts Over Time

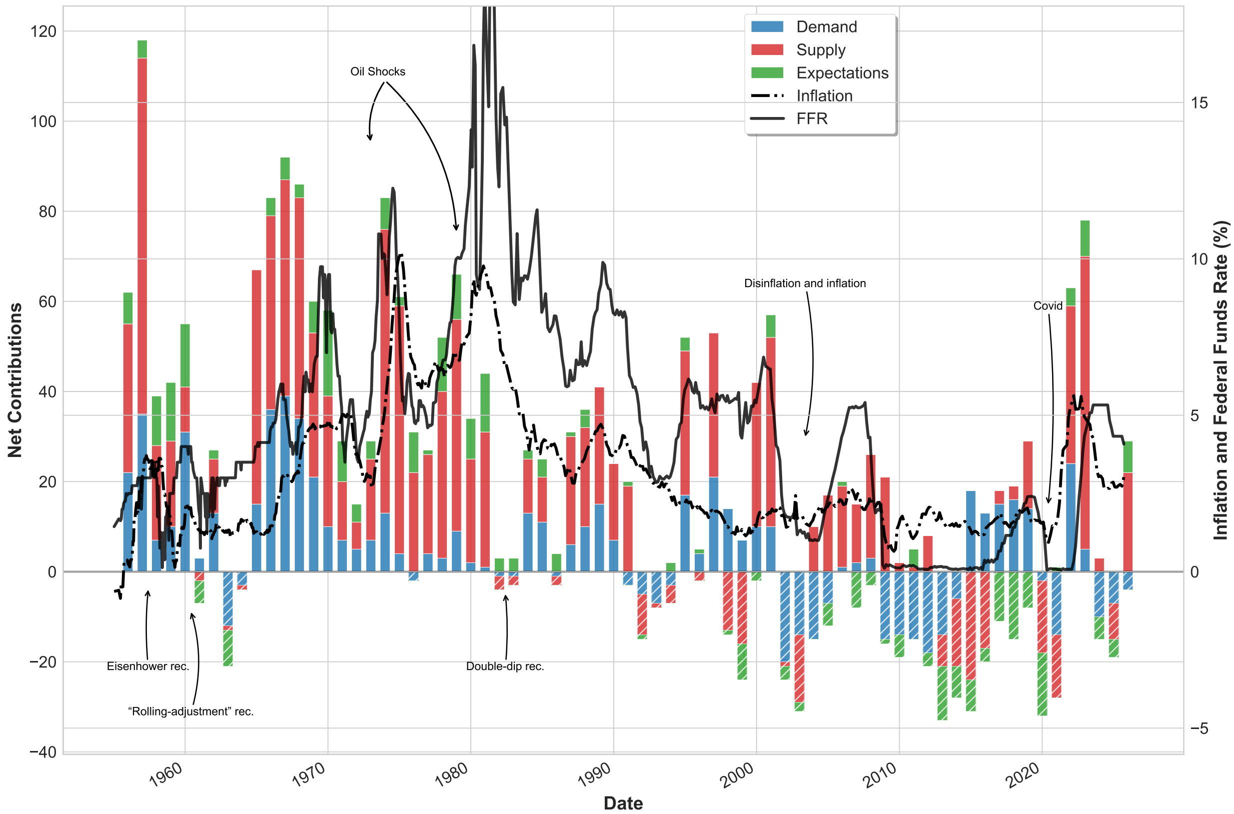

Notes: The figure plots annual net reason counts for demand, supply, and expectations on the left axis. Net imbalance equals the number of reasons indicating higher inflation minus those indicating lower inflation in a given year and category. The right axis shows the Federal Funds Rate and inflation. We use CPI inflation before 1960 and core inflation thereafter. See the main text for variable construction details.

"Rolling-adjustment" recession (1960 – 61). Late 1959 inflation fears were driven by rapid inventory accumulation, strong car sales, and larger defense orders. A sharp turn to inventory liquidation and unemployment above 6 percent flipped the narrative: the Committee now cited slack demand and balance-of-payments worries as disinflationary. Recovery began in mid-1961 as liquidation slowed and defense spending rose, yet persistent labor slack kept inflation subdued. The figure shows many inflationary mentions in 1960, then a marked shift toward disinflationary reasons through 1962.

Oil Shocks (1973, 1979). Throughout the 1970s, we see a relative increase in the presence of supply reasons discussed by the Committee. The 1973 and 1974 meetings mention “oil supply curtailment”, the “fuel situation”, “oil situation”, “increased prices of commodities”, “oil shortages” and “imported oil costs” as supply reasons contributing to an increase of inflation. Similarly, the 1979 meetings mention “higher prices of motor fuels”, “energy price increases”, “energy costs”, while the 1980 meetings already account for “energy price moderation”.

Double-Dip Recession (1980, 1981 - 82). The graph shows mostly inflationary reasons in 1980, followed by a sharp rise in disinflationary reasons over the next years. In early 1980, the Committee blamed the resurgence of prices on “energy costs,′′ “wage growth′′ that stayed at the “rapid pace′′ of 1979, a “depreciating dollar”, and a “widening foreign trade deficit′′. By March, however, the minutes pointed to “rising U.S. interest rates′′ and a “new program of fiscal, energy, credit, and other measures′′ that would push inflation down. Spring and summer entries discuss “weakening demand—declining retail sales,′′ “rising unemployment,′′ and “slow growth of money supply′′—while also noting “energy price moderation.′′ As the slump deepened in 1981-82 the discussion focused on a “strong dollar,′′ “slower wage growth,′′ “persistently high unemployment,′′ and repeated “decline in real GNP,′′ factors that sustained the disinflation despite temporary shocks such as the 1980 “surge in food prices.”

Disinflation and Inflation (2000 - 2006). The graph shows a predominance of inflationary reasons in the 1990s, followed by a surge in disinflationary reasons up to 2003, when we see a reversal. At the turn of the decade, policymakers still saw disinflationary contributions from “increasing productivity,” expected to “hold down increases in unit labor costs and prices,” yet warned that “sharp increases in energy prices” and “higher import prices” together with “taut labor markets” were tilting risks upward.

By mid-2000 the list of cost-push worries had broadened to “increasing medical costs” and the “rise in other import prices,” while demand looked overheated thanks to “strong consumer demand” and “extremely tight labor markets.” The picture reversed in March, 2001: minutes stressed demand shortfalls—“declining consumer confidence,” “weak business investment,” “high inventory levels,” and “low resource utilization”—plus supply-side relief from “falling energy prices” and “highly competitive conditions in most product markets.” Indeed, Bernanke (2003) documents the Fed’s concern with disinflation at this point and its shift toward more aggressive easing.

As the recovery firmed, the narrative flipped again: from 2004 on officials cited "energy prices had risen steeply," "unit labor costs had moved notably higher," "tight resource utilization," and a "recent decline in the dollar" that risked "putting upward pressure on import prices." By 2006 the Committee worried that "elevated energy prices'' together with "continuing increases in resource utilization" and signs that "inflation expectations had risen somewhat'' could entrench the rebound in inflation.

Covid and After (2020 - 2025). Early in the crisis the Committee stressed that "weaker demand and significantly lower oil prices are holding down consumer price inflation," reflecting concern about both demand- and supply-driven disinflation. By mid-2021 the story flipped: the minutes note mostly cost-push pressures, such as "supply chain bottlenecks", "supply constraints", "input cost pressure", as the main contributors to inflation. In 2022, officials acknowledged that price pressures had become broad-based, mentioning "demand-supply imbalances", "higher food and energy prices," and "Russia's war against Ukraine" as contributing to inflation. By 2023, and after the interest rate liftoff, perceived demand and supply contributions fell, while the expectations component rose. This pattern is consistent with Reis (2022*a*), who attributes the 2021–2022 surge to misdiagnosed shocks (overstated slack and persistent supply constraints labeled transitory) and to a late recognition that expectations were drifting from target under a prolonged accommodative stance.

3 Model and Asset Pricing Implications

This section presents a parsimonious New Keynesian framework that links perceived demand and supply disturbances to inflation, monetary policy, and asset prices. Derivations appear in Appendix A.

Identical inflation readings can reflect different underlying shocks, so inflation alone is not a sufficient statistic for policy or asset pricing. In some meetings, yields move mainly through shifts in the expected short-rate path; in other meetings, they move largely through term premia, with long bonds behaving like risky assets and comoving with equities. Policy responses also vary: in some episodes the Committee tightens sharply, while in others inflation remains high and the policy path adjusts less than what the Taylor rule implies. These patterns fit poorly with a unidimensional view of inflation, but arise naturally when inflation is decomposed into different sources.

Demand disturbances correspond to the divine coincidence case: optimal policy offsets them one-for-one, so their pricing channel operates through the expected short-rate path. Supply disturbances generate a stabilization trade-off — inflation rises while real activity weakens — inducing a positive covariance between nominal-bond payoffs and the pricing kernel. The model maps the sign and magnitude of each component into (i) the expected short-rate path, (ii) the term premium, and (iii) the equity risk premium.

Macro Block. Let $x_t := y_t - z_t$ denote the output gap, π_t inflation, and i_t the nominal policy rate, all in log deviations from long-run values. The economy is described by the New Keynesian IS and Phillips curves:

$$x_t = -\psi(i_t - \mathbb{E}_t \pi_{t+1}) + \mathbb{E}_t x_{t+1} + \kappa_d d_t, \tag{1}$$

$$\pi_t = \lambda x_t + \beta \mathbb{E}_t \pi_{t+1} + \phi_s s_t, \tag{2}$$

where d_t and s_t are demand and supply shocks with autoregressive laws of motion

$$d_t = \rho_d d_{t-1} + \varepsilon_{d,t}, \qquad s_t = \rho_s s_{t-1} + \varepsilon_{s,t}, \tag{3}$$

with $\rho_d,\ \rho_s \in (0,1)$, $\kappa_d,\ \phi_s \in (0, +\infty)$, and innovations $\varepsilon_{d,t}$ and $\varepsilon_{s,t}$ independent, mean zero, and conditionally normal. The policy instrument is the short nominal interest rate (i_t) and expectations are common knowledge.

Under discretionary optimal policy with quadratic loss, the equilibrium policy rule is linear in the perceived components:

$$i_t = \mathbb{E}_t \pi_{t+1} + \Upsilon_s s_t + \Upsilon_d d_t, \tag{4}$$

where Υ_s, Υ_d depend only on structural parameters. Demand disturbances are offset one-for-one in equilibrium, but supply disturbances generate a stabilization trade-off. Under standard parameter values,

$$0 < \Upsilon_s < \Upsilon_d, \tag{5}$$

so the real rates respond more to demand than to supply pressures.

Implication 1 (Policy Asymmetry). *The policy rate responds more to demand than to supply disturbances. Supply shocks are only partially offset because of the output-inflation trade-off.*

Asset Prices. Under optimal policy, inflation and the output gap depend only on perceived supply:

$$\pi_t = k^s_\pi s_t, \qquad x_t = k^s_x s_t, \tag{6}$$

with

$$k^s_\pi = \frac{\alpha\phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)} > 0, \qquad k^s_x = -\frac{\lambda\phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)} < 0. \tag{7}$$

Demand disturbances are fully neutralized and do not appear in Equation (6).

Implication 2 (Inflation Attribution). *Under optimal policy, all variation in realized inflation is attributable to the supply component: $\pi_t = k^s_\pi s_t$.*

Aggregate consumption equals output, $C_t = Y_t$, $c_t = \log C_t$. With CRRA preferences, the log SDF is

$$m_{t+1} = \log\beta - \gamma\,(c_{t+1} - c_t). \tag{8}$$

Using $y_t = x_t + z_t$ and the CGG interpretation that supply shocks directly lower output and potential output by the same amount, let $z_t = \bar{z} - z_s s_t$ (with $z_s > 0$). Output therefore falls with supply shocks: $y_t = \bar{z} + y_s s_t$ with $y_s < 0$. Hence,

$$m_{t+1} = \tilde{m}_t - \gamma\, y_s s_{t+1}, \tag{9}$$

where $\tilde{m}_t$ is a term known at time t.

We start by investigating the model implications for bonds. For the two-period nominal yield, the log-linear decomposition following Cieslak and Pflueger (2023) is

$$i_t^{(2)} = \frac{1}{2}(i_t + \mathbb{E}_t i_{t+1}) + \frac{1}{2}\mathrm{Cov}_t\left(m_{t+1}, \pi^e_{t+1}\right) - \left[\frac{1}{2}\mathrm{Cov}_t\left(\pi_{t+1}, \pi^e_{t+1}\right) + \frac{1}{4}\mathrm{Var}_t\left(\pi^e_{t+1}\right)\right], \tag{10}$$

where $\pi^e_{t+1} := \mathbb{E}_{t+1}\pi_{t+2}$. The first term is the risk-neutral expectations component; the covariance term is the inflation risk premium; the last is a Jensen's inequality correction.

Using Equation (4) and the autoregressive dynamics for the demand and supply shocks, the risk-neutral component is linear in (s_t, d_t):

$$\frac{1}{2}(i_t + \mathbb{E}_t i_{t+1}) = \theta_s s_t + \theta_d d_t, \tag{11}$$

and under standard calibrations $\theta_d > \theta_s$.

Implication 3 (Expected Rate Path). *Demand shocks move the risk-neutral component of yields more than supply shocks.*

Because $\pi^e_{t+1} = k^s_\pi \rho_s s_{t+1}$ from (6), combining with (9) yields

$$\mathrm{Cov}_t\left(m_{t+1}, \pi^e_{t+1}\right) = -\gamma\, y_s\, k^s_\pi\, \rho_s\, \mathrm{Var}_t(\varepsilon_{s,t+1}) \;>\; 0, \tag{12}$$

since $y_s < 0$. Supply shocks make nominal bonds risky: they raise rates endogenously, so bond prices fall precisely when consumption is low. Bonds therefore carry a positive risk premium, like equities. By contrast, demand-driven inflation carries no risk premium because policy fully offsets it.

We assume the conditional variance of future supply shocks is increasing in the realization of current supply shocks:

$$\mathrm{Var}_t(\varepsilon_{s,t+h}) \;=\; \mu_s + \tau^h_s s_t, \qquad \tau^h_s > 0, \quad h = 1, 2, \ldots. \tag{13}$$

This reduced-form assumption is made for tractability. New Keynesian models with habit formation generate time-varying risk premia through a different channel: risk aversion rises endogenously when consumption falls close to habit, so that homoskedastic shocks carry larger premia in bad states (Campbell, Pflueger, and Viceira, 2020; Pflueger and Rinaldi, 2022; Pflueger, 2025). These models introduce additional state variables and parameters that are not essential to the demand-supply identification we emphasize.

Implication 4 (Inflation Risk Premium). *Supply disturbances generate a positive inflation risk premium; demand disturbances generate none. The term premium in the yield decomposition loads exclusively on supply volatility.*

Finally, we consider the model implications for the equity market. We follow Pflueger and Rinaldi (2022) and let $P^c_{1,t} = \mathbb{E}_t[M_{t+1}C_{t+1}]$ denote the price of a one-period consumption claim. A log-linear approximation implies

$$\mathbb{E}_t[r^c_{t+1} - r_t] + \frac{1}{2}\text{Var}_t(r^c_{t+1}) = -\text{Cov}_t(m_{t+1}, r^c_{t+1}) = \gamma y_s^2 \text{Var}_t(s_{t+1}) > 0. \tag{14}$$

Because output and consumption fall when s_t rises, the right-hand side is positive when supply shocks are present. Without supply shocks, consumption is fully stabilized and the equity premium is zero.

Implication 5 (Equity Risk Premium). *Supply disturbances generate a positive equity risk premium proportional to* $\text{Var}_t(s_{t+1})$*. Demand disturbances contribute nothing to the consumption-claim premium.*

The next section brings these predictions to the data.

4 The Fed's diagnosis of inflation and monetary policy

In this section, we construct meeting-level demand and supply imbalances and use them to investigate the model's implications. Although supply imbalances track realized inflation more strongly than demand imbalances, an augmented Taylor rule reveals a policy asymmetry: rates respond to demand-driven inflation but partially look through supply pressures, especially from Volcker through the Great Financial Crisis. Yield-curve decompositions show that demand pressures raise risk-neutral yields, whereas supply pressures raise term premia. Pre-meeting macroeconomic conditions explain only a small share of the Fed's diagnosis, and the imbalance measures predict yield changes beyond those macro characteristics.

4.1 Empirical measure for the Fed's diagnosis

The previous section shows that our methodology broadly captures the economic discussions in some of the key historic monetary policy episodes. Next, we use the reasons identified previously to construct measures of the demand and supply proportional contributions to inflation.[9]

As Figure IA.2 showed, reason counts differ across meetings, and a single meeting can cite both inflationary and disinflationary forces within the same category. To address these two facts, we normalize the within-date counts and define the category-x imbalance as

$$I_{x,t} := \frac{N_{x,t}^{(+)} - N_{x,t}^{(-)}}{\sum_{z \in C} \left(N_{z,t}^{(+)} + N_{z,t}^{(-)} \right)}. \tag{15}$$

where d and s denote demand and supply, and C is the set of all categories (demand, supply, expectations, and monetary policy). To limit noise from minor mentions, we restrict the sample to reasons with importance scores of at least five (on a zero to ten scale).[10]

Imbalances lie in $[-1, 1]$. A positive value means upward forces in category x dominate downward ones; a negative value means the reverse. Normalizing by total mentions across all categories yields a relative measure that accounts for the breadth of discussion at each meeting. Internet Appendix IA.1 considers two alternative normalizations: applying the inverse hyperbolic sine to the numerator of Equation (15), and replacing its denominator with an eight-meeting moving average. Results are robust to these choices.

[9] Our main analysis uses the 764 regularly scheduled FOMC meetings for which meeting minutes exist and the LLM extracted at least one inflation-related reason. We exclude two unscheduled meetings: October 6, 1979 (Volcker's emergency meeting) and March 15, 2020 (the Covid emergency rate cut). Intermeeting conference calls never produced standalone minutes and are excluded by construction. Including the two unscheduled meetings produces only borderline significance shifts, concentrated in Post-GFC Taylor rule coefficients. The main results are unaffected.

[10] This restriction drops about 20% of the reasons. Results are not sensitive to this threshold.

4.2 Imbalances and inflation

In the model of Section 3, both demand and supply innovations ($\varepsilon_{d,t}$ and $\varepsilon_{s,t}$) exert pressure to change inflation, although demand pressures are fully offset by the policy response in equilibrium. The imbalance measures $I_{d,t}$ and $I_{s,t}$ are the empirical counterparts of these innovations: they capture the net direction and intensity of demand and supply pressures as perceived by the FOMC at each meeting. We first test whether the imbalances covary with realized inflation by estimating:

$$\pi_t = \alpha_0 + \beta_d I_{d,t} + \beta_s I_{s,t} + \rho \pi_{t-2} + u_t, \tag{16}$$

where π_t is core or headline inflation. Because inflation is highly persistent, we include a lagged dependent variable; we use a two-meeting lag (π_{t-2}) to approximate the quarterly frequency standard in the literature, a convention we maintain throughout. Conditional on π_{t-2}, the coefficients β_d and β_s capture the association between the imbalances and the change in inflation over the two intervening meetings.

Columns 1 and 2 of Table 2 confirm this. Coefficients are larger for headline than for core, consistent with core stripping out food and energy. The supply coefficient is twice as large, consistent with supply factors driving more of the variation in inflation.[11]

We next examine how the imbalances relate to real activity. The dependent variables are Greenbook current-quarter nowcasts of unemployment and GDP growth, and contemporaneous log industrial production:

$$y_t = \alpha_0 + \beta_d I_{d,t} + \beta_s I_{s,t} + \rho y_{t-2} + u_t, \tag{17}$$

where y_t is the activity measure. Columns 3–5 report the results. Demand imbalances enter with the expected signs: negative for unemployment, positive for GDP growth and industrial production. Supply imbalances are insignificant for all three activity measures.[12]

[11] Internet Appendix Tables IA.1 and IA.4 report results under the alternative normalizations in Internet Appendix IA.1. Results are similar across specifications.

[12] Appendix B.2 benchmarks our textual imbalances against the PCE-based demand-supply decomposition of Shapiro (2024). For core PCE, each imbalance loads exclusively on its PCE counterpart. For headline, which unlike core includes energy categories, I_s also loads on demand-driven inflation. This is consistent with Shapiro's (2024, Figure 7) finding that oil supply shocks raise the demand-driven contribution to headline PCE through cross-substitution into energy substitutes. See Appendix B.2 for details.

Table 2: Regressions of inflation and real activity on imbalances

	Inflation		*Activity*		
	Core	*Headline*	*Unemp.*	*GDP Gr.*	*log IP*
I_d	0.099**	0.201***	-0.496***	0.963**	0.015**
	(0.043)	(0.074)	(0.127)	(0.434)	(0.006)
I_s	0.212***	0.392***	0.023	-0.115	0.002
	(0.051)	(0.084)	(0.059)	(0.339)	(0.003)
π_{t-2}	0.980***	0.961***			
	(0.017)	(0.020)			
y_{t-2}			0.958***	0.681***	0.991***
			(0.017)	(0.057)	(0.008)
Obs.	636	636	425	425	762
R^2	0.979	0.956	0.955	0.495	0.995

Notes: Columns 1–2: the dependent variable is the year-over-year level of core or headline PCE inflation from FRED. Columns 3–5: the dependent variables are the Greenbook unemployment rate nowcast, Greenbook real GDP growth nowcast, and log industrial production. I_d and I_s are demand and supply imbalances (Equation (15)). All regressions include a two-meeting lag of the dependent variable (π_{t-2} or y_{t-2}). Newey–West standard errors with 8 lags in parentheses. *, **, *** denote significance at the 10, 5, and 1 percent levels.

4.3 Augmented Taylor rule

The standard Taylor rule raises the policy rate more than one-for-one with inflation regardless of its source. We augment the rule to let demand and supply imbalances enter separately.

The standard Taylor rule with interest rate smoothing is:

$$i_t = \beta_0 + \rho i_{t-2} + (1-\rho)\left[\beta_\pi \pi_t + \varphi_x x_t\right] + \varepsilon_t. \tag{18}$$

Interest-rate smoothing is a well-documented feature of monetary policy, though the underlying mechanism remains debated.[13] As before, i_t is the interest rate, π_t is inflation, x_t

[13] A number of papers have been written on this topic, and a comprehensive overview is outside the scope of this text. An early account by Goodfriend (1991) documents the preference of the Fed for interest rate smoothing. Sack and Wieland (2000) argue that interest rate smoothing might be optimal with forward-looking expectations or uncertainty about the data and the parameters of the model. Coibion and Gorodnichenko (2012) use forecasts from the Greenbook and the SPF and finds evidence that the interest rate

is the log output gap, and β_0 is a reduced-form intercept.[14] During zero lower bound periods, we replace i_t with the shadow rate of Wu and Xia (2016) (available 1990–2022) and use the federal funds rate elsewhere. We use the policy rate lagged two meetings (i_{t-2}) in the right-hand side of the equation above to approximate a quarterly lag used by most of the literature.

We define a Taylor rule that includes our demand and supply components:[15]

$$i_t = \beta_0 + \rho i_{t-2} + (1-\rho)\left[\varphi_\pi \pi_t + \varphi_x x_t + \varphi_d I_{d,t} + \varphi_s I_{s,t}\right] + \varepsilon_t. \tag{19}$$

Following Carvalho, Nechio, and Tristão (2021), we estimate the structural parameters of Equations (18) and (19) by OLS in two steps. First, we define the auxiliary reduced-form equation:

$$i_t = \beta_0^{aux} + \rho i_{t-2} + \varphi_\pi^{aux} \pi_t + \varphi_x^{aux} x_t + \varphi_d^{aux} I_{d,t} + \varphi_s^{aux} I_{s,t} + \varepsilon_t. \tag{20}$$

We estimate Equation (20) by OLS and recover the structural coefficients as

$$\widehat{\varphi}_{\text{var}} = \frac{\widehat{\varphi}_{\text{var}}^{\text{aux}}}{1-\widehat{\rho}}, \quad \text{for} \quad \text{var} = \pi, x, s, d.$$

Standard errors are computed using the delta method.

4.4 Taylor rule estimates under different monetary policy regimes

Clarida, Galí, and Gertler (2000) distinguish two postwar regimes: a pre-Volcker era (1951–1979) in which the Fed largely accommodated expected inflation, and a Volcker–Greenspan era (1979–2006) in which policy responded aggressively.[16] Our demand-supply decomposition lets us revisit this debate and measure how the Fed's response varied across these

inertia likely comes from smoothing, rather than from persistent shocks.

[14] Sack and Wieland (2000) and others make a distinction between this type of policy rule and the original Taylor rule presented by Taylor (1993). We don't make such a distinction, and for the purposes of our analysis, refer to all these variations of interest rate policy responses as "Taylor rules".

[15] Hofmann, Manea, and Mojon (2024) perform a similar exercise, augmenting the baseline specification with demand and supply components of inflation using the decomposition from Shapiro (2024).

[16] For the early U.S. experience, see Friedman and Schwartz (1963); for recent surveys, see Blinder (2022) and Bernanke (2022).

regimes. Table 3 re-estimates our augmented Taylor rule specification by subsample. In all specifications, the dependent variable is the shadow rate whenever it is available (1990-2022), and the federal funds rate otherwise.[17] We use core inflation lagged one month for π_t and contemporaneous log output gap for x_t.

Table 3: Augmented Taylor Rule Estimates

	Pre-Volcker	*Volcker–GFC*	*Post-GFC*
φ_d	-0.71	3.57***	1.88
	(1.12)	(0.85)	(2.50)
φ_s	0.78	-0.19	-0.36
	(0.82)	(0.70)	(2.32)
φ_π	0.76**	1.86***	2.68**
	(0.31)	(0.17)	(1.29)
φ_x	1.11***	0.65***	0.59*
	(0.39)	(0.17)	(0.31)
ρ	0.89***	0.6***	0.91***
	(0.04)	(0.08)	(0.04)
N	269	228	140
R^2	0.917	0.900	0.962

Notes: Each row reports coefficients from the augmented generalized Taylor rule in Equation (19), for different samples: *Pre-Volcker* = 1960Q1–1979Q2; *Volcker-GFC* = 1979Q3–2007Q4; *Post-GFC* = 2008Q1–2025Q2. The dependent variable is the shadow rate from Wu and Xia (2016) whenever that rate is available (1990-2022) and the federal funds rate otherwise. φ_π is the response to π_t (core inflation, lagged one month), φ_x to x_t (log output gap); φ_d and φ_s to the demand and supply textual imbalances. ρ is the interest-rate smoothing parameter. Coefficients are estimated from Eq. (20) and transformed as in Carvalho, Nechio, and Tristão (2021). We construct Newey-West standard errors (with 8 lags) from the auxiliary OLS regressions, and then apply the delta method to generate standard errors for the structural parameters. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

The first column covers the pre-Volcker era—Chairs Martin (1951–70), Burns (1970–78), and Miller (1978–79).[18] The baseline coefficient on inflation is below unity, consistent with the accommodative stance documented by Clarida, Galí, and Gertler (2000). The imbalance coefficients are small and insignificant, suggesting the absence of a targeted policy response during this period.

[17] The results are qualitatively the same when we use the federal funds rate at all times.

[18] The regression begins in 1960, when the FRED's core inflation series becomes available.

The Volcker–GFC column yields sharply different estimates. The inflation coefficient rises to 1.86, well above unity, confirming a proactive anti-inflation stance. The demand-imbalance coefficient is 3.57 and highly significant, while the supply coefficient is statistically indistinguishable from zero. The Fed tightened aggressively when it attributed inflation to demand but did not respond to perceived supply pressures beyond their effect on inflation.

The last column restricts the sample to the post-GFC period (2008–2025). The zero lower bound was binding for much of this period; consequently, persistence dominates the interest rate dynamics in this sample, and most estimates are not statistically significant. Nevertheless, the point estimate for demand remains positive, while that for supply is negative and smaller in magnitude. More importantly, as we show in the next section, imbalances still strongly co-move with the yield curve during this period, linking the Fed's perceptions of inflation's sources to market expectations about the future path of interest rates.[19]

Robustness. Supply pressures affect both contemporaneous inflation and output, so including these controls may attenuate $\hat{\varphi}_s$. We therefore re-estimate the rule excluding π_t and x_t. Table D.3 reports the results. The demand-imbalance coefficient exceeds the supply coefficient in all samples except the pre-Volcker period, consistent with a less differentiated policy stance before Volcker.

A related concern is that regime differences reflect changes in data vintages rather than policy behavior. Following Nakamura, Riblier, and Steinsson (2025), we re-estimate the augmented Taylor rule with two modifications. First, we replace π_t with annualized, real-time, one-quarter-ahead expected inflation, constructed from the Greenbook through 2019:Q4 and from the Survey of Professional Forecasters thereafter. Second, we use a real-time output gap series from the Greenbook through 2019:Q4 and from the CBO thereafter. Appendix Table D.4 reports the results. The estimated differences across periods persist and, if anything, widen when we use real-time gaps and expected inflation (Post-GFC coefficients remain imprecisely estimated throughout).

[19]Appendix B.1 repeats this exercise separately with reasons presented by FOMC members and staff. Internet Appendix Tables IA.2 and IA.5 report results under the alternative normalizations described in Internet Appendix IA.1. Results are very similar across these specifications. Internet Appendix IA.2 includes expectation imbalances in the estimated policy response.

In sum, the Fed moved from a passive, undifferentiated response before 1979 to an active rule under Volcker–Greenspan that leaned heavily against demand-driven inflation while partially looking through supply. After the GFC, unconventional tools muted the policy rate's responsiveness to both inflation and its perceived sources.

Realized inflation and demand-supply pressures. If the Fed leans more aggressively against demand-driven inflation, demand imbalances should have shorter-lived effects on inflation than supply imbalances, whose inflationary consequences the Fed partially accommodates. We re-estimate Equation (16) with block-averaged imbalances—the mean of contemporaneous through three-meeting lags and the mean of four–seven-meeting lags —and control for inflation four and eight meetings earlier (about half a year and one year). Table 4 shows that inflation's dependence on imbalances fades faster for demand than for supply.

A natural question is whether this asymmetry reflects different persistence of the underlying shocks. One of the common explanations for weaker responses to supply is that supply shocks are less persistent, so partially looking through supply-driven inflation is less costly. To test this, we estimate separate autoregressive models, using again averages of adjacent lags for each imbalance. Specifically, we regress contemporaneous imbalance on three subsequent averages: averages of lags one to three, averages of lags four to six, and averages of lags seven to nine. Appendix Figure C.2 plots the estimated coefficients. Both imbalances exhibit similar persistence in our sample; demand and supply are equally persistent in the Fed's discussions about the causes of inflation.

4.5 Yield curve responses to the Fed's inflation diagnosis

The theoretical effect of demand and supply shocks on bond yields is ambiguous. Pflueger (2025) shows that, in a model with external habits, the sign of the nominal-yield response to a demand shock depends on the calibration. When monetary policy reacts aggressively to supply shocks, short rates comove more with output, increasing the covariance of bond returns with consumption. Bonds become more "stock-like," and investors demand a larger premium.

In this section, we examine how the Fed's assessment of the sources of inflation is re-

Table 4: Response of inflation to lagged imbalances

	Dependent variable: π_t			
	Core		*Headline*	
	(1)	(2)	(3)	(4)
$I_{d,t}$ to $I_{d,t-3}$	0.134	-0.042	0.237	-0.034
	(0.171)	(0.380)	(0.272)	(0.560)
$I_{s,t}$ to $I_{s,t-3}$	0.776***	0.803***	1.329***	1.648***
	(0.175)	(0.268)	(0.250)	(0.354)
$I_{d,t-4}$ to $I_{d,t-7}$		-0.100		-0.392
		(0.301)		(0.396)
$I_{s,t-4}$ to $I_{s,t-7}$		1.142***		1.468***
		(0.302)		(0.383)
π_{t-4}	0.939***		0.890***	
	(0.031)		(0.036)	
π_{t-8}		0.839***		0.745***
		(0.052)		(0.059)
Observations	634	630	634	630
R^2	0.946	0.857	0.902	0.785

Notes: The table reports OLS regressions of inflation π_t on demand and supply imbalances at the FOMC-meeting frequency. For each imbalance ($I_{d,t}, I_{s,t}$), we use block-averaged lags: the mean of contemporaneous through 3-meeting lags and, where included, the mean of 4–7 meeting lags. Columns (1) and (2) use core inflation, while columns (3) and (4) use headline inflation. Columns (1) and (3) include only t to $t-3$ for the imbalances and control for four lags of inflation (π_{t-1}–π_{t-4}), capturing short-run persistence. Columns (2) and (4) include both t to $t-3$ and $t-4$ to $t-7$ for the imbalances and instead control for π_{t-8} (roughly one year) to absorb longer-run persistence. Newey–West standard errors with 8 lags are reported in parentheses. See the main text for more details on the construction of the imbalances.

flected in Treasury yields and term premia. Movements in the yield curve reflect both the market's expectations of the future path of short rates and the compensation investors require for bearing interest-rate risk. We use the decomposition of Adrian, Crump, and Moench (2013), which separates Treasury yields into expectations and term-premium components.[20] Let $y_t^{(n)}$ and $tp_t^{(n)}$ denote, respectively, the fitted zero-coupon yield and the term premium at maturity n, measured two business days after FOMC meeting t. The

[20] Data from the Federal Reserve Bank of New York's website.

corresponding risk-neutral yield $rn_t^{(n)}$ is then given by

$$rn_t^{(n)} = y_t^{(n)} - tp_t^{(n)}. \tag{21}$$

We regress risk-neutral yields and term premia at one- and five-year maturities, measured two days after the meeting, on the imbalances, controlling for core inflation, the log output gap, and two-meeting lags of each component, following Equation (20). Table 5 reports results for the post-Volcker period (August 1979 – June 2025). The left panel shows the results for the risk-neutral yields. The demand imbalance is positively related to yields; by contrast, the effect of supply imbalances is much smaller and insignificant. The right panel presents the effects on term premia and reveals another asymmetry. Demand imbalances are negatively and not significantly related to term premia, whereas supply imbalances are positively and significantly related to term premia along the curve.[21]

Panel A of Appendix Table D.5 reports the same regressions for the post-GFC period (2007–2025). The asymmetry persists: demand imbalances load on risk-neutral yields, supply imbalances load on term premia, with larger and more significant supply coefficients in this subsample. The relationship between the Fed's diagnosis and yield-curve movements thus survives the ZLB period, even though the federal funds rate was constrained.[22] Panel B of Appendix Table D.5 reports the pre-Volcker subsample (August 1961 – July 1979). No coefficients are significant, consistent with the pre-Volcker Taylor rule results in Table 3.

The results for one-year yields are consistent with the predictions of the model. Demand-driven inflation triggers a stronger policy tightening path, leading to an increase in risk-neutral yields with little effect on term premia. By contrast, supply imbalances create a

[21] Internet Appendix Table IA.11 re-estimates these regressions using risk-neutral yields and term premia measured fifteen business days after the meeting, an interval that includes the release of the minutes. Demand imbalance coefficients on risk-neutral yields increase in magnitude and significance, and their negative effect on term premia, insignificant at $t+2$, becomes significant. Supply imbalance effects on term premia are unchanged.

[22] A concern is that movements in long yields may largely reflect repricing at short horizons. As a check, we repeat the analysis using Treasury forward rates implied by the Gürkaynak, Sack, and Wright (2007) yield-curve estimates (available from the Federal Reserve Board website). These series are not decomposed into risk-neutral and term-premium components, but they allow us to trace horizon-by-horizon movements in the term structure. We find that both demand and supply imbalances load significantly on forward-rate changes, with demand effects concentrated at shorter horizons and supply effects relatively more pronounced at longer horizons, which is consistent with the results above.

Table 5: Movements in risk-neutral yields and term premia following the Fed's diagnosis

	Risk-Neutral Yields		*Term Premia*	
	1y	*5y*	*1y*	*5y*
I_d	0.480**	0.470***	-0.086	-0.133
	(0.227)	(0.181)	(0.056)	(0.083)
I_s	0.095	0.121	0.115***	0.157**
	(0.106)	(0.075)	(0.036)	(0.061)
π_t	0.269***	0.194***	-0.003	-0.011
	(0.071)	(0.041)	(0.015)	(0.018)
x_t	0.144***	0.103***	-0.015*	-0.037**
	(0.038)	(0.029)	(0.009)	(0.017)
$rn_{t-2}^{(n)}$	0.797***	0.817***	0.019*	0.042**
	(0.043)	(0.033)	(0.010)	(0.017)
$tp_{t-2}^{(n)}$	0.557***	0.152***	0.775***	0.876***
	(0.178)	(0.058)	(0.062)	(0.033)
Obs.	367	367	367	367
R^2	0.951	0.954	0.803	0.908

Notes: Each column reports OLS estimates of risk-neutral yields $rn_t^{(n)}$ or term premia $tp_t^{(n)}$ on demand and supply imbalances (I_d, I_s), contemporaneous core inflation (π_t) and log output gap (x_t), and the risk-neutral yields and term premia lagged by two meetings, analogous to the auxiliary regression in Equation (20). Maturities are one and five years. Risk-neutral yields and term premia come from Adrian, Crump, and Moench (2013); yields and premia are measured two business days after FOMC meeting t. The sample spans August, 1979 to June, 2025 at the FOMC-meeting frequency. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

tradeoff between stabilizing inflation and supporting output. The resulting policy uncertainty raises term premia while exerting a smaller effect on risk-neutral yields.

That said, our baseline model predicts a more limited impact of demand and supply pressures on long-term yields. The persistence we document for longer-maturity yields is instead consistent with models in which markets learn about the Federal Reserve's reaction function by observing policymakers' assessments of demand and supply pressures. For example, Bauer, Pflueger, and Sunderam (2024) estimate time-varying perceptions of the policy rule from Blue Chip survey data; these perceptions update after policy actions and are inversely related to term premia.[23]

[23] Relatedly, Stein and Sunderam (2018) shows that the Federal Reserve's private information about its

How demand and supply attributions interact with market learning about the Fed's reaction function remains an open question for future work. In particular, it would be useful to study how markets combine the Fed's assessment of demand and supply pressures with their beliefs about the Fed's sensitivity to inflation and its policy stance.

4.6 How predictable are the Fed's discussions?

Previous literature such as Cieslak (2018) and Bauer and Swanson (2023*a*,*b*) has shown that changes in interest rate futures around FOMC announcements are partially predictable from past macroeconomic data. We study to what extent the imbalances we extract from FOMC records are similarly predictable. As predictors, we consider a set of pre-announcement macro and financial variables from Bauer and Swanson (2023*b*): the latest nonfarm payroll surprise, year-over-year nonfarm payroll employment growth, three-month S&P 500 returns, the three-month change in the yield-curve slope, three-month changes in commodity prices (BCOM), and the implied skewness of the ten-year Treasury yield.[24] We start the analysis in 1988 due to the availability of the macro news variables.

Columns (1)–(2) of Table 6 regress the demand and supply imbalances on these macro predictors, core inflation, the output gap, and the lagged policy rate. S&P 500 returns positively predict demand imbalances and negatively predict supply imbalances. Commodity price changes load strongly on the supply imbalance. Nonfarm payroll growth and the payroll surprise are positively related to both measures, but significantly only for demand. The R^2 is approximately 30% for each, indicating that pre-meeting macro conditions explain a material share of the variation in imbalances but leave the majority unexplained.

target rate generates policy gradualism and affects long-term rates through the market's inference problem.

[24]The nonfarm payroll surprise (NFP Surprise) is the difference between the most recent nonfarm payroll release and the median expectation for that release from the Money Market Services survey. We standardize the surprise in the NFP to make it similar to the units of the other controls. Employment growth (NFP) is the log change in nonfarm payroll employment from one year earlier to the most recent FOMC announcement. S&P 500 returns are computed over the three months preceding the announcement. Yield Curve Slope is the change in the slope from three months prior to the announcement to one day before the announcement. Commodity prices (BCOM) are the log change in the Bloomberg Commodity Spot Price Index from three months before the announcement to the day before the announcement. Treasury Skewness is the implied skewness of the ten-year Treasury yield over the preceding month. See Bauer and Swanson (2023*b*) for more details. The data is obtained from the San Francisco Fed's website.

Table 6: Imbalances and policy responses controlling for macro news

	Imbalances		*FFR/ Shadow Rate*	
	I_d (1)	I_s (2)	*Volcker–GFC* *(3)*	*Post-GFC* *(4)*
I_d			0.387***	0.008
			(0.095)	(0.221)
I_s			0.017	0.251
			(0.074)	(0.165)
S&P 500	0.438**	-0.766**	-0.262	0.942
	(0.189)	(0.320)	(0.552)	(1.000)
BCOM	0.214	1.505***	1.586***	-1.305*
	(0.196)	(0.329)	(0.446)	(0.765)
NFP Surp.	0.019**	0.027*	0.222**	-0.017
	(0.008)	(0.016)	(0.112)	(0.017)
NFP	0.035**	0.013	0.094**	0.018
	(0.015)	(0.015)	(0.047)	(0.018)
Slope	-0.029	0.019	-0.543***	-0.394***
	(0.035)	(0.049)	(0.075)	(0.148)
Skew	-0.043	0.223**	0.408***	-0.258
	(0.069)	(0.099)	(0.139)	(0.156)
π_t	-0.015	0.025	0.162***	0.186**
	(0.019)	(0.027)	(0.038)	(0.076)
x_t	0.036**	0.049**	0.129***	0.027
	(0.016)	(0.022)	(0.039)	(0.041)
i_{t-2}	-0.007	0.018*	0.855***	0.931***
	(0.009)	(0.009)	(0.024)	(0.041)
Obs.	285	285	158	127
R^2	0.289	0.307	0.982	0.959

Notes: The first two columns report the regressions of the demand and supply imbalances on the expanded set of macro news predictors from Bauer and Swanson (2023*b*) augmented with the log output gap (x_t), core inflation (π_t), and the policy rate lagged two meetings (i_{t-2}). The last two columns re-estimate regressions of the policy rate (i_t) on demand and supply imbalances, with the same set of controls. The policy rate is the shadow rate from Wu and Xia (2016) whenever that rate is available (1990–2022) and the federal funds rate otherwise. The Volcker–GFC period is between February 1988 and December 2007. We start the analysis in 1988 due to the availability of the macro news variables. The post-GFC period is between January 2008 and December 2023. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Columns (3)–(4) regress the policy rate on the imbalances with the same controls, estimated separately for the Volcker–GFC (1988–2007) and post-GFC (2008–2023) periods. In the Volcker–GFC sample, demand imbalances load positively and significantly on the policy rate; supply imbalances do not. The implied long-run demand coefficient is 2.6, exceeding unity and broadly consistent with the estimates in Table 3. Adding macro news controls leaves the core asymmetry intact and, if anything, strengthens it. In the post-GFC period, neither imbalance is significant.

A natural question is whether the imbalances contain information about yield-curve movements beyond what pre-meeting macro conditions already capture. Table 7 re-estimates the yield-curve regressions adding the macro news and the same baseline controls. The asymmetry survives: demand imbalances remain significantly related to risk-neutral yields, while supply imbalances remain significantly related to term premia.[25]

Table 7: Response of risk-neutral yields and term premia to imbalances, controlling for macro news

	Risk-Neutral Yields		*Term Premia*	
	1y	5y	1y	5y
I_d	0.385***	0.365***	-0.015	0.038
	(0.117)	(0.122)	(0.037)	(0.040)
I_s	0.092	0.050	0.059**	0.089***
	(0.056)	(0.061)	(0.025)	(0.029)
Obs.	285	285	285	285
R^2	0.979	0.965	0.883	0.971
Controls	✓	✓	✓	✓

Notes: The table re-estimates regressions of risk-neutral yields ($rn_t^{(n)}$) and term premia ($tp_t^{(n)}$) on demand and supply imbalances, with the expanded set of macro news predictors of Bauer and Swanson (2023*b*) augmented with the log output gap (x_t), core inflation (π_t), and the risk-neutral rate and term premia lagged two meetings. Risk-neutral yields and term premia come from the Adrian, Crump, and Moench (2013) decomposition. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

[25] Internet Appendix Tables IA.3 and IA.6 report results under the alternative normalizations described in Internet Appendix IA.1. Results are very similar across these specifications. Additionally, Internet Appendix Table IA.12 presents the results for the post-GFC subsample.

Daily changes in the yield curve. Appendix Table D.6 reports OLS regressions of three-day changes (one day before to two days after) in risk-neutral yields and term premia on demand and supply imbalances, with and without the macro-news controls. The asymmetry persists: term-premium changes load positively on supply imbalances, while risk-neutral yield changes load positively on demand imbalances (insignificantly so at longer maturities). The coefficients are smaller than in Table 7, consistent with markets incorporating the Fed's inflation assessment gradually rather than within the narrow event window.

4.7 Monetary Policy Shocks

The preceding results show that imbalances predict the level of risk-neutral yields and term premia between FOMC meetings. We now ask whether they also predict asset price changes around each announcement.

High-Frequency Asset Price Changes. A large literature constructs monetary policy surprises from high-frequency changes in federal funds futures, Eurodollar futures, and other asset prices in narrow windows around FOMC announcements.[26] Because minutes and transcripts are released with a lag, market participants observe only the post-meeting statement in real time. To match the information set available at announcement time, we apply our methodology to FOMC statements. The procedure follows Section 2; Internet Appendix IA.7.4 provides the prompt and additional details.

Internet Appendix Table IA.13 confirms that statement demand (supply) imbalances load positively and significantly on minutes demand (supply) imbalances. The R^2 values reveal two further patterns: the mapping from deliberations to statements is noisy, and the R^2 for demand imbalances is roughly twice that for supply, suggesting that demand

[26] Kuttner (2001) uses a scaled version of the current-month federal funds futures (FF1) as his monetary policy shock. Nakamura and Steinsson (2018) and Gürkaynak, Sack, and Swanson (2005) use scaled versions of FF1 and next-month federal funds futures (FF2), along with Eurodollar futures contracts expiring in two to four quarters (ED2, ED3, and ED4). Swanson (2021) includes S&P 500 futures, FF1, FF2, ED1 to ED4, and Treasury yields with maturities of 2, 5, 10, and 30 years, among other series. Although these studies differ in the exact set of series chosen and the methods used to aggregate them into factors, most of them rely on high-frequency changes in asset prices around FOMC announcements. See Brennan et al. (2024) for a comparison of the data and methods used in this literature.

considerations from the meeting are more likely than supply to appear in the public statement.

We consider five series: changes in S&P 500 futures, current-quarter and three-quarter-ahead Eurodollar futures, VIXY, and the monetary policy shock of Bauer and Swanson (2023*b*), all measured in a window from ten minutes before to twenty minutes after the FOMC announcement.[27] We regress each series on statement-based demand and supply imbalances starting in May 1999.[28] Table 8 reports these regressions, both with and without the macro-news controls from Section 4.6.

Table 8: High-Frequency Monetary Policy Shock Regressions

	MPS		*S&P 500*		*ED1*		*ED4*		*VIXY*	
I_d^{stat}	0.008	0.001	-0.040	-0.016	0.006	0.001	0.007	-0.001	0.001	-0.002
	(0.007)	(0.006)	(0.060)	(0.063)	(0.004)	(0.004)	(0.010)	(0.008)	(0.003)	(0.003)
I_s^{stat}	0.028***	0.021*	-0.243***	-0.227***	0.023***	0.020**	0.027**	0.018	0.011***	0.011***
	(0.010)	(0.012)	(0.092)	(0.083)	(0.009)	(0.010)	(0.011)	(0.013)	(0.003)	(0.002)
Obs.	184	184	184	184	184	184	184	184	95	95
R^2	0.053	0.189	0.035	0.058	0.060	0.129	0.031	0.181	0.081	0.148
Controls	✗	✓	✗	✓	✗	✓	✗	✓	✗	✓

Notes: The table reports regressions of high-frequency changes in asset prices and the monetary policy shock of Bauer and Swanson (2023*b*) on the demand and supply imbalances extracted from FOMC statements. Asset price changes are from ten minutes before to twenty minutes after the announcement. The dependent variables include the change in the S&P 500 e-mini futures, current-quarter and three-quarter-ahead Eurodollar futures (ED1 and ED4), VIXY, and the monetary policy shock from Bauer and Swanson (2023*b*). VIXY is an exchange-traded fund tracking short-term VIX futures; changes in VIXY are computed from TAQ data. Regressions without controls use the same set of meetings as those with controls. The control set corresponds to the macro news predictors in Bauer and Swanson (2023*b*), described in Section 4.6. Sample runs from May 1999 to December 2023. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Supply imbalances have a positive and significant effect on Eurodollar futures and the MPS, predict higher volatility, and lower equity futures. Demand imbalances are insignificant in Table 8 despite loading on the policy rate and risk-neutral yields in earlier

[27]The changes in the S&P 500, Eurodollar contracts, and the monetary policy shock are from Bauer and Swanson (2023*b*) and available on the San Francisco Fed website. We use the S&P 500 e-mini futures contract. VIXY is an ETF of short-term VIX futures, available from 2011 on; VIXY changes are constructed manually from TAQ. Starting in January 2023, ED1 and ED4 are the changes for the SOFR futures contracts.

[28]After May 1999, the Fed began releasing a statement after every meeting, regardless of whether policy changed; see Meade and Acosta (2015) for a discussion of the evolution of statement release practices.

specifications. High-frequency changes around announcements capture only the surprise component of the statement, and this asymmetry is consistent with demand assessments being easier for markets to anticipate and appearing more reliably in the statement than supply assessments.

These results align with the model's predictions. Supply-driven inflation confronts the policymaker with a tradeoff between stabilizing inflation and supporting output. Because the Fed cannot fully stabilize both output and inflation under supply pressures, supply imbalances generate risk premia in equity and bond markets. When inflation stems from demand, the policy response stabilizes both, and asset-price reactions are muted.

News articles. Financial markets form their own views on the likely policy stance. We ask whether media coverage of upcoming FOMC meetings captures the public's perception of the Fed's inflation diagnosis, and whether statement-based imbalances predict asset-price movements beyond information already embedded in pre-meeting news coverage.

We collect news articles that discuss the next FOMC meeting and apply our procedure to extract counts of reasons for higher and lower inflation in the same categories as before. We then construct normalized demand and supply imbalance measures, $I_{d,t}^{news}$ and $I_{s,t}^{news}$. Internet Appendix IA.7.5 details the article collection, filtering, and processing, and reports the prompts used to extract reasons. As a validation exercise, Internet Appendix Table IA.13 regresses statement-based imbalances on news-based imbalances and confirms that statement demand (supply) imbalances load only on news demand (supply) imbalances. A similar pattern holds for minutes.

The differences $I_{d,t}^{stat} - I_{d,t}^{news}$ and $I_{s,t}^{stat} - I_{s,t}^{news}$ capture gaps between the Fed's assessment of demand and supply contributions to inflation and the public's diagnosis. Internet Appendix Table IA.14 regresses the same asset price changes on $I_{d,t}^{stat} - I_{d,t}^{news}$ and $I_{s,t}^{stat} - I_{s,t}^{news}$, controlling for the macro predictors. The supply-gap coefficient remains significant: a wider divergence between the Fed's and the public's supply assessment predicts lower equity prices and higher volatility.[29]

[29]Internet Appendix Table IA.15 reports an alternative specification that regresses the same high-frequency asset price changes on statement imbalances, controlling for news imbalances and the macro predictors; results are qualitatively unchanged. Internet Appendix Table IA.16 adds news imbalances to the yield-curve regressions of Table 7. The effect of supply imbalances on term premia is unchanged. News

Media coverage tracks, but does not subsume, the Fed's inflation diagnosis. News-based demand (supply) imbalances predict statement demand (supply) imbalances, yet the residual gap, especially for supply, helps explain high-frequency asset-price responses. The statements therefore convey information about the Fed's inflation diagnosis beyond what pre-meeting news articles already reflect.

Our stylized model replicates the main empirical patterns. Supply pressures raise risk premia on stocks and bonds, whereas demand pressures do not. By contrast, the expected path of future rates responds more to demand than to supply pressures, reflecting the greater sensitivity of the policy rate to demand in the baseline CGG model.

5 Conclusion

We use large language models to extract the FOMC's inflation attributions from the full historical record of meeting minutes. The resulting panel of classified statements yields three main results.

First, realized inflation rises more after supply-attributed episodes than demand-attributed ones, as standard New Keynesian models predict when monetary policy leans more aggressively against demand than supply. The classification thus tracks economically meaningful variation in underlying shocks.

Second, augmented Taylor rules show the funds rate responds more strongly to demand-driven inflation. The asymmetry is absent pre-Volcker and strengthens afterward, consistent with a post-1980 regime of pre-empting demand overheating while partly looking through supply pressures.

Third, higher perceived demand raises risk-neutral yields, while higher perceived supply raises the bond and equity premia. Demand pressures involve no stabilization trade-off and raise expected short rates. Supply pressures create a trade-off and increase interest-rate uncertainty and term premia. These results survive controls for pre-meeting economic conditions, indicating that the Fed's real-time assessment carries predictive content beyond what past data alone would imply.

These findings carry important implications. First, the Fed's diagnosis of inflation

controls attenuate the demand coefficient on risk-neutral yields, consistent with demand assessments being more predictable by the market.

contains incremental predictive content for both macro outcomes and policy actions. Researchers and market participants should therefore track not just the magnitude of inflation but also its perceived origin. Second, policy-rule estimates that ignore the demand–supply distinction risk conflating heterogeneous episodes and understating systematic variation in the reaction function. Third, the LLM-based classification extends naturally to other dimensions of central-bank communication, such as financial-stability concerns or distributional objectives, without sacrificing transparency or scalability.

What policymakers say about inflation's origins, not just its level, shapes both their actions and market outcomes. The tools to extract and quantify these narratives at scale now exist, opening the door to questions that were previously out of reach.

References

Adrian, Tobias, Richard K. Crump, and Emanuel Moench. 2013. "Pricing the term structure with linear regressions." *Journal of Financial Economics*, 110(1): 110–138.

Aruoba, S. Borağan and Thomas Drechsel. 2024. "Identifying Monetary Policy Shocks: A Natural Language Approach." National Bureau of Economic Research Working Paper 32417.

Baqaee, David and Emmanuel Farhi. 2022. "Supply and demand in disaggregated Keynesian economies with an application to the Covid-19 crisis." *American Economic Review*, 112(5): 1397–1436.

Bauer, Michael D. and Eric T. Swanson. 2023*a*. "An Alternative Explanation for the 'Fed Information Effect'." *American Economic Review*, 113(3): 664–700.

Bauer, Michael D. and Eric T. Swanson. 2023*b*. "A Reassessment of Monetary Policy Surprises and High-Frequency Identification." *NBER Macroeconomics Annual*, 37: 87–155.

Bauer, Michael D., Carolin E. Pflueger, and Adi Sunderam. 2024. "Perceptions about Monetary Policy." *Quarterly Journal of Economics*, 139(4).

Beaudry, Paul, Thomas J Carter, and Amartya Lahiri. 2023. "The Central Bank's Dilemma: Look Through Supply Shocks or Control Inflation Expectations?" National Bureau of Economic Research Working Paper 31741.

Bennani, Hamza, Tobias Kranz, and Matthias Neuenkirch. 2018. "Disagreement between FOMC members and the Fed's staff: New insights based on a counterfactual interest rate." *Journal of Macroeconomics*, 58: 139–153.

Bernanke, Ben. 2022. *21st Century Monetary Policy.* W. W. Norton & Company, Incorporated.

Bernanke, Ben S. 2003. "An Unwelcome Fall in Inflation?" *Remarks before the Economics Roundtable, University of California, San Diego, La Jolla, California*.

Black, Fischer. 1995. "Interest Rates as Options." *The Journal of Finance*, 50(5): 1371–1376.

Blinder, Alan S. 2022. *A monetary and fiscal history of the United States, 1961-2021.* Princeton University Press.

Brennan, Connor M., Margaret M. Jacobson, Christian Matthes, and Todd B. Walker. 2024. "Monetary Policy Shocks: Data or Methods?" *Finance and Economics Discussion Series*, , (2024–011): 1–49.

Campbell, John Y., Carolin Pflueger, and Luis M. Viceira. 2020. "Macroeconomic Drivers of Bond and Equity Risks." *Journal of Political Economy*, 128(8): 3148–3185.

Carvalho, Carlos, Fernanda Nechio, and Tiago Tristão. 2021. "Taylor rule estimation by OLS." *Journal of Monetary Economics*, 124: 140–154.

Cieslak, Anna. 2018. "Short-Rate Expectations and Unexpected Returns in Treasury Bonds." *The Review of Financial Studies*, 31(9): 3265–3306.

Cieslak, Anna and Carolin Pflueger. 2023. "Inflation and Asset Returns." *Annual Review of Financial Economics*, 15(1): 433–448.

Cieslak, Anna and Michael McMahon. 2023. "Tough Talk: The Fed and the Risk Premium." Working paper.

Cieslak, Anna, Stephen Hansen, Michael McMahon, and Song Xiao. 2023. "Policymakers' uncertainty." National Bureau of Economic Research.

Clarida, Richard, Jordi Galí, and Mark Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *The Quarterly Journal of Economics*, 115(1): 147–180.

Cogley, Timothy and Thomas J Sargent. 2001. "Evolving post-world war II US inflation dynamics." *NBER macroeconomics annual*, 16: 331–373.

Cogley, Timothy and Thomas J Sargent. 2005. "Drifts and volatilities: monetary policies and outcomes in the post WWII US." *Review of Economic dynamics*, 8(2): 262–302.

Coibion, Olivier and Yuriy Gorodnichenko. 2012. "Why Are Target Interest Rate Changes so Persistent?" *American Economic Journal: Macroeconomics*, 4(4): 126–162.

Curti, Filippo and Sophia Kazinnik. 2023. "Let's face it: Quantifying the impact of nonverbal communication in FOMC press conferences." *Journal of Monetary Economics*, 139: 110–126.

di Giovanni, Julian, Sebnem Kalemli-Ozcan, Alvaro Silva, and Muhammed Yildirim. 2023. "Global Supply Chain Pressures, Trade, and Inflation." Working paper / conference slides.

Eichengreen, Barry and Peter M. Garber. 1991. "Before the Accord: U.S. Monetary-Financial Policy, 1945-51."

Eickmeier, Sandra and Boris Hofmann. 2022. "What drives inflation? Disentangling demand and supply factors." Bank for International Settlements. Forthcoming in *International Journal of Central Banking*.

Engelberg, Joseph, Asaf Manela, William Mullins, and Luka Vulicevic. 2025. "Entity Neutering."

Friedman, Milton and Anna Jacobson Schwartz. 1963. *A Monetary History of the United States, 1867-1960.* Princeton University Press.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. 2019. "Text as data." *Journal of Economic Literature*, 57(3): 535–74.

Goodfriend, Marvin. 1991. "Interest rates and the conduct of monetary policy." *Carnegie-Rochester Conference Series on Public Policy*, 34: 7–30.

Goodfriend, Marvin and Robert G. King. 2005. "The incredible Volcker disinflation." *Journal of Monetary Economics*, 52(5): 981–1015.

Gorodnichenko, Yuriy. 2021. "Comment on "The People versus the Markets: A Parsimonious Model of Inflation Expectations"." *Brookings Papers on Economic Activity*, Fall: 370–377.

Gorodnichenko, Yuriy, Tho Pham, and Oleksandr Talavera. 2023. "The Voice of Monetary Policy." *American Economic Review*, 113(2): 548–584.

Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Ivan Werning. 2021. "Monetary Policy in Times of Structural Reallocation." *SSRN Electronic Journal*.

Gürkaynak, Refet S, Brian Sack, and Eric Swanson. 2005. "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking*, 1(1).

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright. 2007. "The U.S. Treasury yield curve: 1961 to the present." *Journal of Monetary Economics*, 54(8): 2291–2304.

Hansen, Anne Lundgaard and Sophia Kazinnik. 2023. "Can ChatGPT Decipher Fedspeak?" *SSRN Electronic Journal*.

Hansen, Stephen, Michael McMahon, and Andrea Prat. 2017. "Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach." *The Quarterly Journal of Economics*, 133(2): 801–870.

Hofmann, Boris, Cristina Manea, and Benoit Mojon. 2024. "Targeted Taylor rules: some evidence and theory." Bank for International Settlements BIS Working Papers 1234.

Kim, Chang-Jin and Charles R Nelson. 2006. "Estimation of a forward-looking monetary policy rule: A time-varying parameter model using ex post data." *Journal of Monetary Economics*, 53(8): 1949–1966.

Kuttner, Kenneth N. 2001. "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market." *Journal of Monetary Economics*, 47(3): 523–544.

Lucca, David and Francesco Trebbi. 2009. *Measuring Central Bank Communication: An Automated Approach with Application to FOMC Statements.*

Meade, Ellen E. and Miguel Acosta. 2015. "Hanging on Every Word: Semantic Analysis of the FOMC's Postmeeting Statement." *FEDS Notes*, 2015.0(1580).

Nakamura, Emi and Jón Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *The Quarterly Journal of Economics*, 133(3): 1283–1330.

Nakamura, Emi, Venance Riblier, and Jón Steinsson. 2025. "Beyond the Taylor Rule." Working Paper.

Orphanides, Athanasios. 2002. "Monetary-Policy Rules and the Great Inflation." *American Economic Review*, 92(2): 115–120.

Orphanides, Athanasios. 2003. "Historical monetary policy analysis and the Taylor rule." *Journal of Monetary Economics*, 50(5): 983–1022.

Pflueger, Carolin. 2025. "Back to the 1980s or not? The drivers of inflation and real risks in Treasury bonds." *Journal of Financial Economics*, 167: 104027.

Pflueger, Carolin and Gianluca Rinaldi. 2022. "Why does the Fed move markets so much? A model of monetary policy and time-varying risk aversion." *Journal of Financial Economics*, 146(1): 71–89.

Reis, Ricardo. 2022*a*. "The burst of high inflation in 2021–22: how and why did we get here?" Bank for International Settlements BIS Working Papers 1060.

Reis, Ricardo. 2022*b*. "Losing the Inflation Anchor." *Brookings Papers on Economic Activity*, 2021(2): 307–379.

Romer, Christina and David Romer. 1989. "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz." National Bureau of Economic Research, Inc NBER Working Papers 2966.

Romer, Christina D and David H Romer. 2002. "A Rehabilitation of Monetary Policy in the 1950's." *American Economic Review*, 92(2): 121–127.

Romer, Christina D and David H Romer. 2008. "The FOMC versus the Staff: Where Can Monetary Policymakers Add Value?" *American Economic Review*, 98(2): 230–235.

Romer, Christina D. and David H. Romer. 2023. "Presidential Address: Does Monetary Policy Matter? The Narrative Approach after 35 Years." *American Economic Review*, 113(6): 1395–1423.

Rotemberg, Julio J. 2013. "Penitence after Accusations of Error: 100 Years of Monetary Policy at the U.S. Federal Reserve." National Bureau of Economic Research. Cambridge, MA. Prepared for the NBER Symposium on "The First Hundred Years of the Federal Reserve," July 10, 2013.

Sack, Brian and Volker Wieland. 2000. "Interest-rate smoothing and optimal monetary policy: a review of recent empirical evidence." *Journal of Economics and Business*, 52(1–2): 205–228.

Shapiro, Adam Hale. 2024. "Decomposing Supply- and Demand-Driven Inflation." *Journal of Money, Credit and Banking*.

Shapiro, Adam Hale and Daniel J. Wilson. 2022. "Taking the Fed at its Word: A New Approach to Estimating Central Bank Objectives using Text Analysis." *The Review of Economic Studies*, 89(5): 2768–2805.

Smets, Frank and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review*, 97(3): 586–606.

Stein, Jeremy C. and Adi Sunderam. 2018. "The Fed, the Bond Market, and Gradualism in Monetary Policy." *The Journal of Finance*, 73(3): 1015–1060.

Svensson, Lars E.O. 1999. "Inflation targeting as a monetary policy rule." *Journal of Monetary Economics*, 43(3): 607–654.

Swanson, Eric T. 2021. "Measuring the effects of federal reserve forward guidance and asset purchases on financial markets." *Journal of Monetary Economics*, 118: 32–53.

Taylor, John B. 1993. "Discretion versus policy rules in practice." *Carnegie-Rochester Conference Series on Public Policy*, 39: 195–214.

Wei, Jason, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2022. "Chain-of-Thought Prompting Elicits Reasoning in Large Language Models."

Wu, Jing Cynthia and Fan Dora Xia. 2016. "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound." *Journal of Money, Credit and Banking*, 48(2–3): 253–291.

Appendix

- Appendix A contains the details of the model.
- Appendix B contains robustness and extensions.
- Appendix C contains supplementary figures.
- Appendix D contains supplementary tables.

A Model

This appendix derives the equilibrium policy rule and asset-pricing implications of the model introduced in Section 3. We restate the key equations here for reference. The IS and New Keynesian Phillips Curve equations are:

$$x_t = -\psi(i_t - \mathbb{E}_t \pi_{t+1}) + \mathbb{E}_t x_{t+1} + \kappa_d d_t \tag{IS}$$

$$\pi_t = \lambda x_t + \beta \mathbb{E}_t \pi_{t+1} + \phi_s s_t \tag{PC}$$

The shock processes follow:

$$d_t = \rho_d d_{t-1} + \varepsilon_{d,t} \tag{22}$$

$$s_t = \rho_s s_{t-1} + \varepsilon_{s,t} \tag{23}$$

with $\varepsilon_{d,t}$, $\varepsilon_{s,t}$ independent, zero-mean, conditionally normal innovations and ρ_d, $\rho_s \in (0,1)$.

The Fed minimizes a quadratic loss over its dual mandate:

$$\min_{i_t} \frac{1}{2} \mathbb{E}_t \left(\sum_{j=0}^{\infty} \beta^j (\alpha x_{t+j}^2 + \pi_{t+j}^2) \right) \tag{$\star$}$$

Under discretion without commitment, the Fed reoptimizes each period, choosing $\{x_t, \pi_t, i_t\}$ subject to (PC) and (IS). Following CGG, the problem decomposes into two steps: first minimize (⋆) over (x_t, π_t) given (PC), then back out i_t from (IS). Since the Fed cannot manipulate private-sector beliefs under discretion, the Lagrangian is:

$$\mathcal{L} = -\frac{1}{2}\left[\alpha x_t^2 + \pi_t^2\right] + \mu(\pi_t - \lambda x_t - f_t) + F_t \tag{24}$$

where μ is the multiplier and

$$f_t = \beta \mathbb{E}_t \pi_{t+1} + \phi_s s_t \tag{25}$$

$$F_t = -\frac{1}{2}\mathbb{E}_t\left[\sum_{j=1}^{\infty} \beta^j (\alpha x_{t+j}^2 + \pi_{t+j}^2)\right] \tag{26}$$

are treated as fixed in the optimization. The FOC yields:

$$x_t = -\frac{\lambda}{\alpha}\pi_t \tag{27}$$

Substituting into the Phillips curve

$$\pi_t = -\frac{\lambda^2}{\alpha}\pi_t + \beta\mathbb{E}_t\pi_{t+1} + \phi_s s_t \tag{28}$$

$$\implies \pi_t = \frac{\alpha}{\alpha+\lambda^2}\beta\mathbb{E}_t\pi_{t+1} + \frac{\alpha}{\alpha+\lambda^2}\phi_s s_t \tag{29}$$

Assuming rational expectations with common beliefs about (d_t, s_t) and the usual transversality conditions:

$$\begin{aligned}\pi_t &= \frac{\alpha}{\alpha+\lambda^2}(\phi_s s_t) + \left(\frac{\alpha}{\alpha+\lambda^2}\right)^2 \beta(\phi_s \rho_s s_t) + \left[\frac{\alpha}{\alpha+\lambda^2}\right]^3 \beta^2(\phi_s \rho_s^2 s_t) + \ldots \\ &= \frac{\alpha}{\alpha+\lambda^2}\phi_s\left[s_t + \frac{\beta\alpha\rho_s}{\alpha+\lambda^2}s_t + \left(\frac{\beta\alpha\rho_s}{\alpha+\lambda^2}\right)^2 s_t + \ldots\right] \\ &= \frac{\alpha\phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)}s_t \end{aligned} \tag{30}$$

It follows that:

$$\mathbb{E}_t \pi_{t+h} = \frac{\alpha\phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)} \rho_s^h s_t, \quad \forall h \geq 0 \tag{31}$$

$$x_t = -\frac{\lambda\phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)} s_t \tag{32}$$

$$\mathbb{E}_t x_{t+h} = -\frac{\lambda\phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)} \rho_s^h s_t, \quad \forall h \geq 0 \tag{33}$$

The demand shock $\kappa_d d_t$ does not appear in equilibrium inflation because the Fed sets interest rates to offset it exactly. Substituting the first-order condition into the IS curve (IS) yields:

$$i_t = \mathbb{E}_t \pi_{t+1} + \underbrace{\frac{1}{\psi}\left[\frac{\lambda(1-\rho_s)\phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)}\right]}_{\Upsilon_s} s_t + \underbrace{\frac{1}{\psi}\kappa_d}_{\Upsilon_d} d_t \tag{34}$$

The real rate $(i_t - \mathbb{E}_t \pi_{t+1})$ fully offsets $\kappa_d d_t$, since $\Upsilon_d = \kappa_d/\psi$ exactly neutralizes the IS disturbance. Supply shocks are only partially offset, with $\Upsilon_s \to 0$ as $\alpha \to \infty$.

To compare magnitudes, suppose $\rho_d = \rho_s = \rho > 0$ and $\kappa_d = \phi_s = 1$. Then $\Upsilon_d > \Upsilon_s$ if and only if:

$$\lambda^2 + \alpha(1-\beta\rho) - \lambda(1-\rho) > 0 \tag{35}$$

A sufficient condition is $\lambda > 1-\rho$, which holds for $\lambda = 0.3$ and $\rho > 0.7$, values consistent with Clarida, Galí, and Gertler (2000). Thus:

$$0 < \Upsilon_s < \Upsilon_d \tag{36}$$

A.1 Asset Prices

Using the equilibrium expressions from Section 3, we derive the asset-pricing results stated there. Recall that under optimal policy, $\pi_t = k_\pi^s s_t$, $x_t = k_x^s s_t$ (Eq. (6)), and $m_{t+1} = \tilde{m}_t - \gamma y_s s_{t+1}$ (Eq. (9)). Since s_{t+1} is conditionally normal, $(m_{t+1}, \pi_{t+1}, \mathbb{E}_{t+1}\pi_{t+2})$ are jointly normal as in Cieslak and Pflueger (2023).

Nominal bonds. Consider nominal risk-free bonds paying one dollar in one and two periods, with prices $\exp(-i_t)$ and $\exp(-2i_t^{(2)})$, and a real one-period bond with price $\exp(-r_t)$. To isolate inflation risk, we treat the real rate r_t as approximately constant. For the two-period nominal yield, we adopt the log-linear decomposition in Cieslak and Pflueger (2023). Under the assumption of an approximately constant real rate,

$$i_t^{(2)} = \frac{1}{2}\big(i_t + \mathbb{E}_t i_{t+1}\big) + \frac{1}{2}\mathrm{Cov}_t\big(m_{t+1}, \pi_{t+1}^e\big) - \left[\frac{1}{2}\mathrm{Cov}_t\big(\pi_{t+1}, \pi_{t+1}^e\big) + \frac{1}{4}\mathrm{Var}_t\big(\pi_{t+1}^e\big)\right], \tag{37}$$

where $\pi_{t+1}^e := \mathbb{E}_{t+1}\pi_{t+2}$. The first term, is the expectations-hypothesis (EH) component: the average expected short rate over the bond's life, equal to the yield under risk-neutrality up to Jensen/convexity adjustments. The second term, is the inflation risk-premium. This covariance term isolates a priced inflation-risk channel that provides a transparent decomposition of the bond risk premium. Following Cieslak and Pflueger (2023), we treat this channel as a tractable benchmark, while recognizing that empirical term premia may also reflect other sources of risk.

Using the equilibrium Taylor rule (Eq. (34)) and the AR(1) processes for (d_t, s_t), the risk-neutral component can be written as

$$\frac{1}{2}\big(i_t + \mathbb{E}_t i_{t+1}\big) = \theta_s s_t + \theta_d d_t, \tag{38}$$

with

$$\theta_s = \frac{1}{2}\left[\frac{\alpha\phi_s\rho_s + \frac{\lambda}{\psi}(1-\rho_s)\phi_s}{\lambda^2 + \alpha(1-\beta\rho_s)}\right](1+\rho_s), \tag{39}$$

$$\theta_d = \frac{1}{2}\frac{\kappa_d}{\psi}(1+\rho_d), \tag{40}$$

Under typical parameter choices (e.g. Clarida, Galí, and Gertler (2000)), we have $0 < \theta_s < \theta_d$. From the equilibrium dynamics $\pi_{t+1}^e = k_\pi^s \rho_s s_{t+1}$ and the SDF $m_{t+1} = \tilde{m}_t - \gamma y_s s_{t+1}$:

$$Cov_t(m_{t+1}, \pi_{t+1}^e) = Cov_t\big(\tilde{m}_t - \gamma y_s s_{t+1}, k_\pi^s \rho_s s_{t+1}\big) = -\gamma y_s k_\pi^s \rho_s Var_t(\varepsilon_{s,t+1}). \tag{41}$$

which is positive since $y_s < 0$, yielding Equation (12).

Risky asset. Following Pflueger and Rinaldi (2022), the risky asset is a claim to future consumption C_{t+h}, $h = 1, \ldots, \infty$, sold at price P^c_t. An aggregate levered firm purchases the consumption claim at P^c_t and finances it with equity worth θP^c_t and one-period risk-free debt worth $(1-\theta)P^c_t$. For simplicity, we focus on a one-period, zero-coupon consumption claim, with price:

$$P^c_{1,t} = \mathbb{E}_t[M_{t+1}C_{t+1}] \tag{42}$$

and its (log) return is approximately

$$r^c_{t+1} - r_t \approx (y_{t+1} - y_t) = y_s(s_{t+1} - s_t).$$

The standard log-linear approximation implies

$$\mathbb{E}_t[r^c_{t+1} - r_t] + \frac{1}{2}\text{Var}_t(r^c_{t+1}) = -\text{Cov}_t(m_{t+1}, r^c_{t+1}). \tag{43}$$

Using $m_{t+1} = \tilde{m}_t - \gamma y_s s_{t+1}$ and r_t constant, we obtain

$$-\text{Cov}_t(m_{t+1}, r^c_{t+1}) = \gamma y_s^2 \text{Var}_t(s_{t+1}) > 0, \tag{44}$$

as stated in Equation (14).

B Robustness and Extensions

This appendix presents two extensions. Appendix B.1 compares the assessments of FOMC members and staff. Appendix B.2 compares our decomposition to a PCE-based alternative by Shapiro (2024).

B.1 Comparing FOMC members and staff

FOMC members do not make decisions in isolation. The Board of Governors' staff gather real-time data, synthesize it with models, and present the economic and financial outlook before each meeting. The relative influence of FOMC members and staff has practical implications. Svensson (1999) separates their roles: staff construct forecasts under alternative policy scenarios, and the FOMC chooses among them. Romer and Romer (2008) compare FOMC and staff forecasts and find that FOMC forecasts add little information beyond staff forecasts, while the forecast difference predicts monetary policy shocks. Bennani, Kranz, and Neuenkirch (2018) document substantial disagreement between the policy rate implied by the staff's Greenbook and the rate chosen by the FOMC.

We use our measure to examine whether staff and FOMC perceptions of the causes of inflation are complementary or whether one group's assessment has a disproportionate impact. As described in Internet Appendix IA.7, we flag each inflation reason as presented by either an FOMC member or staff. We then compute imbalances as in Section 2 separately for each group.

Internet Appendix Figure IA.7 plots the share of reasons discussed by FOMC members over time. The share declines under Burns, then rises to about 80% under Miller and remains high under Volcker and Greenspan. It falls under Bernanke, especially after the GFC, then rises modestly under Yellen and more sharply under Powell.

Table B.1 reports the regressions in Table 2 estimated separately for FOMC members and staff. Coefficients are similar across groups, but the demand coefficient for staff is insignificant. Table B.2 reports the specification in Table 7 estimated separately for FOMC and staff. Coefficients are similar for risk-neutral yields. Term-premium loadings on supply imbalances are larger and statistically significant only for FOMC reasons.

Appendix Table B.3 repeats Taylor rule estimation by monetary policy period separately for imbalances constructed solely from FOMC participants and staff members. The greater responsiveness of the policy rate to demand imbalances during the Volcker-GFC period comes mostly from FOMC members; the policy rule is much less sensitive to staff assessments.

Table B.4 reports horse-race regressions of Greenbook forecasts for inflation and unemployment on demand and supply imbalances for both FOMC and staff, controlling for the dependent variable lagged two periods. Inflation forecasts load more on FOMC supply imbalances than on staff supply imbalances. For both groups, loadings on demand imbalances are much smaller, especially for staff. FOMC demand imbalances are more strongly associated with lower unemployment forecasts, whereas supply imbalances are are much less weakly related for both groups.

The results suggest that both staff and FOMC assessments are relevant and help explain forecasts and policy decisions. However, policy rates load much more strongly on FOMC assessments, indicating that FOMC members' views carry disproportionate weight in the final decision.

B.2 Comparison to the Shapiro (2024) decomposition

Shapiro (2024) decomposes PCE inflation into demand-driven (π_t^D), supply-driven (π_t^S), and ambiguous components by applying sign restrictions to price-quantity comovements at the category level.[30] We benchmark our narrative imbalances against these series.

Our theoretical framework implies that the imbalances relate to changes in inflation, not levels. Consequently, if I_d and I_s capture distinct inflation sources, I_d should load on changes in demand-driven inflation and I_s on changes in supply-driven inflation. We regress Shapiro's components on the narrative imbalances, controlling for the own two-meeting lag, lagged core PCE inflation, and the log output gap:

$$\pi_t^{D/S} = \alpha + \beta_d I_{d,t} + \beta_s I_{s,t} + \rho_D \pi_{t-2}^D + \rho_S \pi_{t-2}^S + \gamma \pi_{t-1} + \delta x_t + u_t, \tag{45}$$

where π_t^D and π_t^S are Shapiro's demand- and supply-driven inflation (year on year), π_{t-1} is lagged core PCE inflation, and x_t is the log output gap.

Table B.5 reports the results. For core PCE (columns 1–2), the two decompositions align: I_d loads positively and significantly on demand-driven inflation but not on supply-driven inflation, and I_s loads on supply-driven inflation but not on demand-driven inflation.

For headline PCE (columns 3–4), I_d enters with the expected positive sign for demand-

[30] The procedure can be summarized as follows. For each PCE category i and month t, run rolling regressions of quantity on lags of quantity and price, and of price on lags of price and quantity; let the one-step residuals be $\hat{u}_{i,t}^Q$ and $\hat{u}_{i,t}^P$. Classify i as "demand" if $\text{sign}(\hat{u}_{i,t}^Q) = \text{sign}(\hat{u}_{i,t}^P)$, "supply" if the signs differ, and "ambiguous" if either residual is near zero. Let the sets of demand and supply categories be denoted by D_t and S_t. Then demand and supply contributions are the expenditure-share-weighted averages of category inflation over the corresponding sets:

$$\pi_t^D = \sum_{i \in D_t} \omega_{i,t-1} \pi_{i,t}, \qquad \pi_t^S = \sum_{i \in S_t} \omega_{i,t-1} \pi_{i,t},$$

and analogously for the ambiguous component.

driven inflation, though the coefficient does not reach conventional significance ($p = 0.13$). I_s loads on both components. Since this cross-loading is present for headline but absent for core, it is likely driven by the food and energy categories that core excludes. An oil price increase, for instance, raises prices and lowers quantities in directly affected products such as fuel oil (supply-driven), but alternative energy products see their relative price fall, inducing substitution, and display co-moving prices and quantities (demand-driven). Shapiro (2024) documents this channel: oil supply shocks raise both the supply-driven and demand-driven contributions to headline inflation through cross-substitution among energy products.

Table B.1: Regressions of inflation on imbalances: FOMC and Staff

	All		*FOMC*		*Staff*	
	Core	*Headline*	*Core*	*Headline*	*Core*	*Headline*
I_d	0.099**	0.201***	0.121***	0.205***	0.035	0.080
	(0.043)	(0.074)	(0.038)	(0.060)	(0.024)	(0.053)
I_s	0.212***	0.392***	0.129***	0.226***	0.129***	0.263***
	(0.051)	(0.084)	(0.035)	(0.065)	(0.035)	(0.059)
π_{t-2}	0.980***	0.961***	0.983***	0.968***	0.985***	0.968***
	(0.017)	(0.020)	(0.019)	(0.023)	(0.018)	(0.021)
Obs.	636	636	636	636	636	636
R^2	0.979	0.956	0.978	0.954	0.978	0.955

Notes: The dependent variables are monthly levels of headline and core inflation, expressed as year-over-year percentage changes, from the FRED. Demand and supply imbalances are defined as in Equation (15). We compute imbalances separately for FOMC members, staff, and the pooled set of reasons. Regressions are estimated by OLS with Newey-West standard errors with 8 lags (in parentheses). *, **, *** denote significance at the 10, 5, and 1 percent levels, respectively.

Table B.2: Response of risk-neutral yields and term premia controlling for macro news – FOMC members and staff

	Risk-Neutral Yields		*Term Premia*	
	1y	*5y*	*1y*	*5y*
Panel A: FOMC Reasons				
I_d	0.197**	0.164**	-0.011	0.013
	(0.088)	(0.075)	(0.025)	(0.033)
I_s	0.069	0.024	0.041**	0.075***
	(0.051)	(0.058)	(0.021)	(0.027)
Obs.	285	285	285	285
R^2	0.978	0.964	0.882	0.971
Controls	✓	✓	✓	✓
Panel B: Staff Reasons				
I_d	0.177***	0.185***	0.012	0.035
	(0.055)	(0.062)	(0.020)	(0.024)
I_s	0.024	0.007	0.005	0.003
	(0.051)	(0.047)	(0.021)	(0.025)
Obs.	285	285	285	285
R^2	0.978	0.965	0.881	0.970
Controls	✓	✓	✓	✓

Notes: The table re-estimates regressions of the risk-neutral yields and term premia on demand and supply imbalances, with the set of macro-news predictors of Bauer and Swanson (2023*b*). We compute the regressions separately for imbalances constructed with FOMC and staff reasons. Risk-neutral yields and term premia come from the Adrian, Crump, and Moench (2013) decomposition. Additionally, each regression controls for the log output gap, core inflation, and the risk-neutral rate and term premia lagged two meetings. See Appendix B.1 for details on the variable definitions. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table B.3: Augmented Generalized Taylor Rule Estimates by Era – FOMC and Staff

	FOMC *Pre-Volcker*	 *Volcker–GFC*	 *Post-GFC*	Staff *Pre-Volcker*	 *Volcker–GFC*	 *Post-GFC*
φ_d	-2.06	2.63***	-0.15	0.01	1.21***	1.96
	(1.40)	(0.51)	(1.38)	(1.48)	(0.29)	(2.14)
φ_s	-0.79	0.17	0.14	1.12	-0.20	-1.40
	(0.84)	(0.52)	(1.48)	(0.90)	(0.53)	(1.68)
φ_π	0.86***	1.90***	2.33**	0.73**	1.87***	2.82*
	(0.30)	(0.19)	(1.13)	(0.31)	(0.19)	(1.44)
φ_x	1.23***	0.70***	0.68**	1.06***	0.68***	0.62*
	(0.43)	(0.19)	(0.29)	(0.37)	(0.21)	(0.32)
ρ	0.89***	0.62***	0.90***	0.89***	0.61***	0.91***
	(0.04)	(0.08)	(0.04)	(0.04)	(0.08)	(0.04)
N	269	228	140	269	228	140
R^2	0.918	0.898	0.962	0.918	0.896	0.962

Notes: Each row reports structural slope coefficients from the augmented generalized Taylor rule in Equation (19), for different samples: *Pre-Volcker* = 1960Q1–1979Q2; *Volcker-GFC* = 1979Q3–2007Q4; *Post-GFC* = 2008Q1–2025Q2. We estimate the structural slopes using imbalances constructed with reasons expressed by FOMC participants and by staff separately. The dependent variable is the shadow rate from Wu and Xia (2016) whenever that rate is available (1990-2022) and the federal funds rate otherwise. φ_π is the response to π_t (core inflation, lagged one month), φ_x to x_t (log output gap); φ_d and φ_s to the demand and supply textual imbalances. ρ is the interest-rate smoothing parameter. Coefficients are estimated from Eq. (20) and transformed as in Carvalho, Nechio, and Tristão (2021). We construct Newey-West standard errors (with 8 lags) from the auxiliary OLS regressions, and then apply the delta method to generate standard errors for the structural parameters. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table B.4: Greenbook Forecasts and FOMC/Staff Imbalances

	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
Panel A: Inflation Greenbook Forecasts					
I_d^{FOMC}	0.077	0.304	0.222*	0.168**	0.124**
	(0.316)	(0.193)	(0.120)	(0.083)	(0.055)
I_s^{FOMC}	1.116***	0.361**	0.149	0.176**	0.151***
	(0.430)	(0.148)	(0.123)	(0.072)	(0.055)
I_d^{Staff}	-0.053	-0.030	-0.021	-0.001	0.030
	(0.261)	(0.115)	(0.077)	(0.049)	(0.043)
I_s^{Staff}	0.332	0.129	0.102	0.131	0.034
	(0.211)	(0.171)	(0.081)	(0.082)	(0.041)
Obs.	278	278	278	278	278
R^2	0.521	0.750	0.911	0.953	0.968

	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
Panel B: Unemployment Greenbook Forecasts					
I_d^{FOMC}	-0.310***	-0.325***	-0.309***	-0.303***	-0.311***
	(0.105)	(0.098)	(0.106)	(0.101)	(0.102)
I_s^{FOMC}	0.020	0.035	0.059	0.061	0.079
	(0.047)	(0.055)	(0.058)	(0.061)	(0.066)
I_d^{Staff}	-0.119**	-0.112*	-0.112	-0.100	-0.099
	(0.052)	(0.062)	(0.070)	(0.071)	(0.073)
I_s^{Staff}	0.037	0.029	0.032	0.026	-0.004
	(0.040)	(0.041)	(0.040)	(0.042)	(0.047)
Obs.	425	410	403	374	344
R^2	0.955	0.948	0.941	0.940	0.943

Notes: The table presents horse-race regressions of Greenbook/Tealbook forecasts for core CPI inflation and unemployment for h quarters ahead against the imbalances for FOMC members and staff. We control for the dependent variable lagged two meetings (roughly one quarter) in all regressions. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table B.5: Regressions of Shapiro (2024) inflation components on imbalances

	Core		*Headline*	
	Demand	*Supply*	*Demand*	*Supply*
I_d	0.104**	-0.075	0.113	-0.011
	(0.047)	(0.052)	(0.070)	(0.089)
I_s	0.034	0.096**	0.166***	0.267***
	(0.028)	(0.042)	(0.049)	(0.066)
π^D_{t-2}	0.477***	-0.333***	0.732***	-0.008
	(0.084)	(0.074)	(0.064)	(0.067)
π^S_{t-2}	-0.323***	0.581***	-0.015	0.878***
	(0.065)	(0.068)	(0.058)	(0.074)
π_{t-1}	0.331***	0.331***	0.084*	0.059
	(0.057)	(0.055)	(0.047)	(0.051)
x_t	0.012	0.017	0.040**	0.034*
	(0.010)	(0.011)	(0.018)	(0.018)
Obs.	478	478	478	478
R^2	0.943	0.977	0.893	0.943

Notes: The dependent variables are Shapiro (2024) demand-driven (π^D_t) and supply-driven (π^S_t) PCE inflation components (year-over-year), for core and headline measures. I_d and I_s are demand and supply narrative imbalances as defined in Equation (15). π^D_{t-2} and π^S_{t-2} are the two-meeting lags of Shapiro demand- and supply-driven inflation (headline or core, matching the column). π_{t-1} is one-month lagged core PCE inflation. x_t is the log output gap. Regressions are estimated by OLS with Newey-West standard errors with 8 lags (in parentheses). *, **, *** denote significance at the 10, 5, and 1 percent levels, respectively.

C Supplementary Figures

Figure C.1: Aggregated Reason Counts Over Time

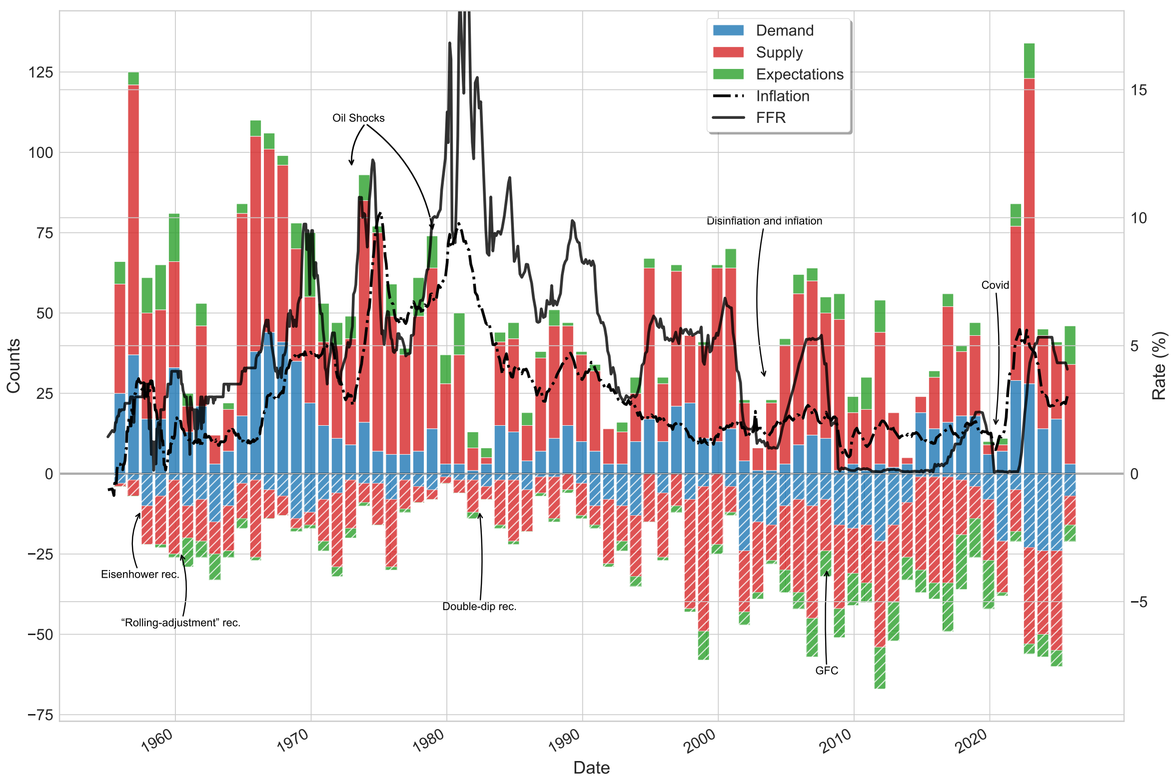

Notes: This figure aggregates the demand, supply, and expectation reason counts annually, separately for reasons increasing (bars going up, full) and decreasing inflation (bars going down, hatched) in the left axis. Additionally, we present the Federal Funds Rate and inflation in the right axis. We use CPI inflation before 1960 and core inflation after that. See the main text for more details on the construction of the demand and supply measures from the minutes.

Figure C.2: Persistence of imbalances

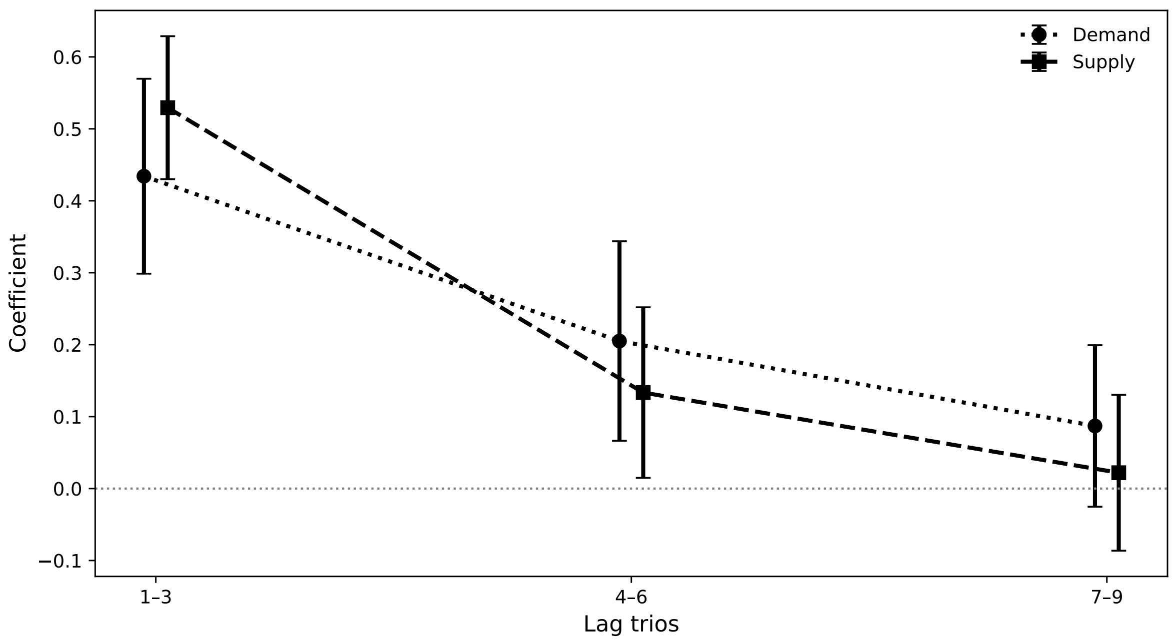

Notes: This figure plots autoregressive coefficients with the associated 95% confidence intervals. Each imbalance is related to the average of adjacent lag trios, estimated separately for demand and supply. Specifically, we regress contemporaneous imbalance on three subsequent averages: averages of lags one to three, averages of lags four to six, and averages of lags seven to nine. We use Newey–West standard errors with 8 lags. See the main text for details of the variable construction.

D Supplementary Tables

Table D.1: Sample Periods and Number of Observations by Series

Series	**Start**	**End**	**Obs.**
FOMC Records			
Full corpus	Jan 1937	Sep 2025	766
Largest estimation sample	Jan 1960	Jun 2025	636
Statements	Aug 1994	Dec 2023	201
News articles	May 1999	Jun 2023	151
Macroeconomic series			
PCE Core Inflation	Jan 1960	Jun 2025	638
PCE Headline Inflation	Jan 1960	Jun 2025	638
Federal Funds Effective Rate	Sep 1954	Sep 2025	725
Shadow Rate	Feb 1990	Jan 2022	256
Monetary policy shock variables			
Macroeconomic News	Feb 1988	Dec 2023	287
Yield series			
Risk-neutral yield	Jun 1961	Sep 2025	618
Term premium	Jun 1961	Sep 2025	618

Notes: This table reports the start date, end date, and total number of observations for each series used in the empirical analysis. See Section 1 for more details about the data sources.

Table D.2: Top reasons for inflation by category

Demand		*Supply*		*Expectations*	
Reason	Count	Reason	Count	Reason	Count
resource slack	17	energy prices	66	stable longer-run inflation expectations	12
high resource utilization	16	dollar depreciation	40	inflation expectations	12
slack in resource utilization	15	productivity gains	25	lower longer-term inflation expectations	12
slack in labor and product markets	14	wage increases	21	stable inflation expectations	9
economic slack	9	food prices	20	rising inflation expectations	8
tight labor markets	8	oil prices	19	stable longer-term inflation expectations	8
labor market strengthening	7	unit labor costs	18	declining inflation expectations	7
underutilized resources (slack)	7	commodity prices	15	well-anchored longer-term inflation expectations	6
tightening resource utilization	7	energy price declines	13	well-anchored inflation expectations	6
tight labor market	6	steel price increases	13	contained inflation expectations	5
Total	**106**	**Total**	**250**	**Total**	**85**

Notes: The table reports the ten most cited explanations for inflation in the minutes, sorted by category. "Demand", "Supply", and "Expectations" reflect the model's tags. "Count" is the number of occurrences across the full sample. See the main text for details on the construction of these variables.

Table D.3: Direct estimates of policy rate and imbalances

	Pre-Volcker	*Volcker-GFC*	*Post-GFC*
$I_{d,t}$	0.156	1.220***	0.540***
	(0.139)	(0.332)	(0.196)
$I_{s,t}$	0.197*	0.242	0.240
	(0.112)	(0.182)	(0.250)
i_{t-2}	0.956***	0.906***	0.977***
	(0.039)	(0.038)	(0.032)
Obs.	269	228	140
R^2	0.907	0.860	0.943

Notes: Each row reports coefficients from the estimate of the policy rate on the policy rate lagged two meetings and demand and supply imbalances. We estimate the regression separately for different samples: *Pre-Volcker* = 1960Q1–1979Q2; *Volcker-GFC* = 1979Q3–2007Q4; *Post-GFC* = 2008Q1–2025Q2. The dependent variable is the shadow rate from Wu and Xia (2016) whenever that rate is available (1990-2022) and the federal funds rate otherwise. We construct Newey-West standard errors (with 8 lags). *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table D.4: Augmented Taylor Rule estimates with real-time data

	Pre-Volcker	*Volcker–GFC*	*Post-GFC*
φ_d	-0.44	4.60***	-0.63
	(1.09)	(1.30)	(2.09)
φ_s	0.96	0.01	-0.06
	(0.87)	(0.81)	(2.17)
φ_π	1.11***	1.69***	2.09
	(0.27)	(0.24)	(1.33)
φ_x	0.53***	0.22*	0.64***
	(0.10)	(0.12)	(0.22)
ρ	0.76***	0.68***	0.90***
	(0.04)	(0.06)	(0.04)
N	157	228	137
R^2	0.874	0.893	0.958

Notes: Each row reports coefficients from the augmented generalized Taylor rule in Equation (19), for different samples: *Pre-Volcker* = 1967Q1–1979Q2; *Volcker-GFC* = 1979Q3–2007Q4; *Post-GFC* = 2008Q1–2025Q2. The dependent variable is the shadow rate from Wu and Xia (2016) whenever that rate is available (1990-2022) and the federal funds rate otherwise. The coefficients are estimated from Eq. (20) and transformed as in Carvalho, Nechio, and Tristão (2021). φ_π measures the response to real-time, one-quarter-ahead, annualized expected inflation. We use the Greenbook-based series through 2019:Q4 and the Survey of Professional Forecasters counterpart thereafter. φ_x measures the response to the real-time output gap constructed from the Greenbook through 2019:Q4 and the real-time CBO series thereafter. φ_d and φ_s measure responses to demand and supply textual imbalances. ρ is the interest-rate smoothing parameter. We construct Newey-West standard errors (with 8 lags) from the auxiliary OLS regressions, and then apply the delta method to generate standard errors for the structural parameters. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table D.5: Movements in risk-neutral yields and term premia following the Fed's diagnosis

	Panel A: Post-GFC				*Panel B: Pre-Volcker*			
	Risk-Neutral		*Term Premium*		*Risk-Neutral*		*Term Premium*	
	1y	*5y*	*1y*	*5y*	*1y*	*5y*	*1y*	*5y*
I_d	0.235*	0.309**	-0.128**	-0.110	-0.066	-0.047	0.006	0.021
	(0.133)	(0.138)	(0.061)	(0.117)	(0.092)	(0.077)	(0.030)	(0.034)
I_s	-0.056	-0.025	0.191***	0.355***	0.035	0.059	-0.009	-0.032
	(0.102)	(0.095)	(0.044)	(0.105)	(0.093)	(0.069)	(0.029)	(0.036)
π_t	0.200**	0.164***	-0.013	-0.064**	0.057	0.057	0.037***	0.020
	(0.077)	(0.044)	(0.014)	(0.028)	(0.048)	(0.042)	(0.012)	(0.015)
x_t	0.059*	0.044	-0.002	-0.004	0.096***	0.066***	-0.005	-0.024***
	(0.035)	(0.027)	(0.011)	(0.023)	(0.027)	(0.019)	(0.007)	(0.008)
$rn_{t-2}^{(n)}$	0.872***	0.835***	-0.012	-0.013	0.881***	0.849***	-0.023	0.028
	(0.048)	(0.056)	(0.010)	(0.031)	(0.048)	(0.057)	(0.015)	(0.024)
$tp_{t-2}^{(n)}$	0.196	-0.013	0.733***	0.839***	0.483*	0.145	0.678***	0.808***
	(0.250)	(0.067)	(0.102)	(0.054)	(0.269)	(0.103)	(0.055)	(0.036)
Obs.	139	139	139	139	247	247	247	247
R^2	0.953	0.936	0.747	0.834	0.911	0.903	0.746	0.897

Notes: Each column reports OLS estimates of risk-neutral yields $rn_t^{(n)}$ or term premia $tp_t^{(n)}$ on demand and supply imbalances (I_d, I_s), contemporaneous core inflation (π_t) and log output gap (x_t), and the $rn_t^{(n)}$ and $tp_t^{(n)}$ lagged by two meetings, analogous to the auxiliary regression in Equation (20). Maturities are one and five years. Risk-neutral yields and term premia come from Adrian, Crump, and Moench (2013); yields and premia are measured two business days after FOMC meeting t. Panel A spans January 2008 to June 2025; Panel B spans August 1961 to July 1979. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table D.6: Imbalances on monetary policy shocks around FOMC announcements

	Δ *Risk-Neutral Yields*				Δ *Term Premia*			
	1y	1y	5y	5y	1y	1y	5y	5y
I_d	0.055**	0.050**	0.043*	0.039	-0.017*	-0.014	0.001	0.007
	(0.022)	(0.023)	(0.023)	(0.027)	(0.010)	(0.012)	(0.016)	(0.019)
I_s	-0.015	-0.024*	-0.007	-0.017	0.021***	0.023***	0.029***	0.031**
	(0.012)	(0.014)	(0.016)	(0.018)	(0.007)	(0.008)	(0.010)	(0.013)
Obs.	287	287	287	287	287	287	287	287
R^2	0.027	0.078	0.012	0.043	0.027	0.031	0.016	0.037
Controls	✗	✓	✗	✓	✗	✓	✗	✓

Notes: The table reports regressions of changes in risk-neutral yields and term premia on supply and demand imbalances, with and without the macro news predictors of Bauer and Swanson (2023*b*). Risk-neutral yields and term premia come from the Adrian, Crump, and Moench (2013) decomposition. Changes are measured as the difference between one day before and two days after each FOMC meeting. Regressions without controls use the same sample as those with controls to facilitate comparison. See the main text for variable definitions. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Internet Appendix

- Internet Appendix IA.1 considers different functional forms.
- Internet Appendix IA.2 adds imbalances from expectations-driven reasons.
- Internet Appendix IA.3 details the neutering exercise as robustness.
- Internet Appendix IA.4 contains supplementary figures.
- Internet Appendix IA.5 contains supplementary tables.
- Internet Appendix IA.6 contains examples of the LLM output.
- Internet Appendix IA.7 contains details of the prompt designs for our LLM methodology.

IA.1 Different Functional Forms

Given a set of categories $X = \{x_1, x_2, x_3, ..., x_N\}$, we define the vector of net differences

$$N_t = \begin{bmatrix} N^{(+)}_{x_1,t} - N^{(-)}_{x_1,t} \\ N^{(+)}_{x_2,t} - N^{(-)}_{x_2,t} \\ \vdots \\ N^{(+)}_{x_N,t} - N^{(-)}_{x_N,t} \end{bmatrix} \tag{IA.1}$$

We consider two alternative normalizations. The first divides the net imbalances by a moving average of the total number of reasons over the last eight meetings,

$$I^{ma}_t := \frac{N_t}{\sum_{j=0}^{7}(\sum_{i=1}^{N} N^{(+)}_{x_i,t-j} + N^{(-)}_{x_i,t-j})/8}. \tag{IA.2}$$

The second applies the inverse hyperbolic sine to the net imbalances,

$$I^{asine}_t := asine(N_t) = \log\left(N_t + \sqrt{N_t^2 + 1}\right). \tag{IA.3}$$

Internet Appendix Tables IA.1, IA.2, and IA.3 replicate Tables 2, 3, and 7 using the moving-average normalization. Internet Appendix Tables IA.4, IA.5, and IA.6 report the same regressions using the inverse hyperbolic sine normalization.

Table IA.1: Regressions of inflation on imbalances (moving average)

	All		FOMC		Staff	
	Core	*Headline*	*Core*	*Headline*	*Core*	*Headline*
I_d^{ma}	0.069	0.107	0.095***	0.120**	0.030	0.050
	(0.044)	(0.070)	(0.033)	(0.054)	(0.026)	(0.042)
I_s^{ma}	0.208***	0.400***	0.107***	0.206***	0.102***	0.178***
	(0.050)	(0.083)	(0.034)	(0.063)	(0.025)	(0.040)
π_{t-2}	0.980***	0.960***	0.983***	0.967***	0.984***	0.968***
	(0.018)	(0.021)	(0.019)	(0.023)	(0.018)	(0.021)
Obs.	636	636	636	636	636	636
R^2	0.979	0.956	0.978	0.954	0.978	0.954

Notes: This table replicates Table 2 using an alternative imbalance specification that scales net counts by a moving average of total counts over the last eight meetings. The dependent variables are monthly headline and core inflation from FRED, expressed as year-over-year percentage changes. Demand and supply imbalances are defined as in Appendix IA.1. Regressions are estimated by OLS with Newey–West standard errors (8 lags, in parentheses). *, **, and *** denote significance at the 10, 5, and 1 percent levels, respectively.

Table IA.2: Augmented Taylor Rule Estimates (moving average)

	Pre-Volcker	*Volcker–GFC*	*Post-GFC*
φ_d	-1.60	4.17***	1.12
	(1.15)	(0.81)	(2.48)
φ_s	1.47	-1.09	-0.96
	(0.99)	(0.67)	(2.34)
φ_π	0.74**	1.89***	2.64**
	(0.30)	(0.16)	(1.19)
φ_x	1.11***	0.69***	0.64**
	(0.37)	(0.15)	(0.31)
ρ	0.89***	0.59***	0.90***
	(0.04)	(0.08)	(0.04)
N	269	228	140
R^2	0.918	0.902	0.962

Notes: This table replicates Table 3 using an alternative imbalance specification that scales net counts by a moving average of total counts over the last eight meetings. Each row reports structural slope coefficients from the augmented generalized Taylor rule in Equation (19), for different samples: *Pre-Volcker* = 1960Q1–1979Q2; *Volcker-GFC* = 1979Q3–2007Q4; *Post-GFC* = 2008Q1–2025Q2. The dependent variable is the shadow rate from Wu and Xia (2016) whenever that rate is available (1990-2022) and the federal funds rate otherwise. φ_π is the response to π_t (core inflation, lagged one month), φ_x to x_t (log output gap); φ_d and φ_s to the demand and supply textual imbalances. ρ is the interest-rate smoothing parameter. Demand and supply imbalances are defined as in Appendix IA.1. Coefficients are estimated from Eq. (20) and transformed as in Carvalho, Nechio, and Tristão (2021). We construct Newey-West standard errors (with 8 lags) from the auxiliary OLS regressions, and then apply the delta method to generate standard errors for the structural parameters. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table IA.3: Response of risk-neutral yields and term premia controlling for macro news (moving average)

	Risk-Neutral Yields		*Term Premia*	
	1y	5y	1y	5y
I_d^{ma}	0.400***	0.398***	-0.004	0.048
	(0.140)	(0.139)	(0.040)	(0.048)
I_s^{ma}	0.084*	0.063	0.065***	0.089***
	(0.051)	(0.059)	(0.024)	(0.028)
Obs.	285	285	285	285
R^2	0.978	0.965	0.884	0.971
Controls	✓	✓	✓	✓

Notes: This table replicates Table 7 using an alternative imbalance specification that scales net counts by a moving average of total counts over the last eight meetings. The table re-estimates regressions of the risk-neutral yields and term premia on demand and supply imbalances, with the set of macro-news predictors of Bauer and Swanson (2023*b*). Demand and supply imbalances are defined as in Appendix IA.1. Risk-neutral yields and term premia come from the Adrian, Crump, and Moench (2013) decomposition. Additionally, each regression controls for the log output gap, core inflation, and risk-neutral yield and term premium lagged two meetings. See the main text for details on the variable definitions. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table IA.4: Regressions of inflation on imbalances (inverse hyperbolic sine)

	All		**FOMC**		**Staff**	
	Core	*Headline*	*Core*	*Headline*	*Core*	*Headline*
I_d^{asine}	0.022*	0.042**	0.035**	0.059**	0.019	0.053*
	(0.012)	(0.020)	(0.014)	(0.024)	(0.019)	(0.029)
I_s^{asine}	0.057***	0.119***	0.043***	0.099***	0.080***	0.150***
	(0.013)	(0.021)	(0.014)	(0.027)	(0.021)	(0.028)
π_{t-2}	0.981***	0.960***	0.984***	0.967***	0.984***	0.966***
	(0.018)	(0.021)	(0.019)	(0.023)	(0.017)	(0.020)
Obs.	636	636	636	636	636	636
R^2	0.979	0.958	0.978	0.955	0.979	0.956

Notes: This table replicates Table 2 using an alternative imbalance specification that applies the inverse hyperbolic sine to the net counts. The dependent variables are monthly headline and core inflation from FRED, expressed as year-over-year percentage changes. Demand and supply imbalances are defined as in Appendix IA.1. Regressions are estimated by OLS with Newey–West standard errors (8 lags, in parentheses). *, **, and *** denote significance at the 10, 5, and 1 percent levels, respectively.

Table IA.5: Augmented Taylor Rule Estimates (inverse hyperbolic sine)

	Pre-Volcker	*Volcker–GFC*	*Post-GFC*
φ_d	-0.15	0.87***	0.43
	(0.38)	(0.13)	(0.46)
φ_s	0.54	-0.14	-0.18
	(0.38)	(0.13)	(0.43)
φ_π	0.70**	1.86***	2.85**
	(0.30)	(0.17)	(1.23)
φ_x	1.02***	0.62***	0.58*
	(0.34)	(0.15)	(0.31)
ρ	0.89***	0.59***	0.91***
	(0.04)	(0.08)	(0.04)
N	269	228	140
R^2	0.918	0.903	0.962

Notes: This table replicates Table 3 using an alternative imbalance specification that applies the inverse hyperbolic sine transformation to net counts. Each row reports structural slope coefficients from the augmented generalized Taylor rule in Equation (19), for different samples: *Pre-Volcker* = 1960Q1–1979Q2; *Volcker-GFC* = 1979Q3–2007Q4; *Post-GFC* = 2008Q1–2025Q2. The dependent variable is the shadow rate from Wu and Xia (2016) whenever that rate is available (1990-2022) and the federal funds rate otherwise. φ_π is the response to π_t (core inflation, lagged one month), φ_x to x_t (log output gap); φ_d and φ_s to the demand and supply textual imbalances. ρ is the interest-rate smoothing parameter. Demand and supply imbalances are defined as in Appendix IA.1. Coefficients are estimated from Eq. (20) and transformed as in Carvalho, Nechio, and Tristão (2021). We construct Newey-West standard errors (with 8 lags) from the auxiliary OLS regressions, and then apply the delta method to generate standard errors for the structural parameters. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table IA.6: Response of risk-neutral yields and term premia controlling for macro news (inverse hyperbolic sine)

	Risk-Neutral Yields		*Term Premia*	
	1y	5y	1y	5y
I_d^{asine}	0.072**	0.074***	0.002	0.011
	(0.031)	(0.028)	(0.007)	(0.009)
I_s^{asine}	0.019	0.009	0.012**	0.021***
	(0.013)	(0.013)	(0.005)	(0.007)
Obs.	285	285	285	285
R^2	0.979	0.965	0.883	0.971
Controls	✓	✓	✓	✓

Notes: This table replicates Table 7 using an alternative imbalance specification that applies the inverse hyperbolic sine to net counts. The table re-estimates regressions of the risk-neutral yields and term premia on demand and supply imbalances, with the set of macro-news predictors of Bauer and Swanson (2023*b*). Demand and supply imbalances are defined as in Appendix IA.1. Risk-neutral yields and term premia come from the Adrian, Crump, and Moench (2013) decomposition. Additionally, each regression controls for the log output gap, core inflation, and risk-neutral yield and term premium lagged two meetings. See the main text for details on the variable definitions. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

IA.2 Including Expectations

We construct a baseline expectations imbalance from the full set of expectation-related reasons identified by the LLM (Internet Appendix IA.7).[31] Internet Appendix Table IA.7 adds the expectations imbalance to the inflation regressions. The expectations coefficient is positive and marginally significant for pooled core inflation, and significant for FOMC members in both core and headline specifications; staff-based expectations imbalances are insignificant.

In the augmented Taylor rule (Internet Appendix Table IA.8), φ_e is positive and significant in the full sample, with the strongest effect in the Volcker–GFC era; the coefficient turns negative and insignificant post-GFC.

Regressions of risk-neutral yields and term premia controlling for macro-news predictors yield insignificant expectations coefficients across all specifications (Internet Appendix Table IA.9). In high-frequency regressions using statement-based imbalances (Internet Appendix Table IA.10), expectations imbalances are associated with lower S&P 500 futures returns and higher volatility, alongside supply pressures.

[31] Inflation expectations are hard to measure, and FOMC language about them evolved over time. We consider two alternatives. A narrower definition reclassifies all reasons originally tagged as Expectations with a stricter prompt requiring explicit reference to long-run expectations (≥5-year horizon) and a documented shift; reasons that no longer qualify revert to their appropriate category. A broader definition reclassifies all reasons originally tagged as Other with a prompt that recognizes historical language such as "inflation psychology" or "market psychology," which the Fed often used in lieu of explicit references to expectations (Gorodnichenko, 2021; Goodfriend and King, 2005; Rotemberg, 2013; Reis, 2022*b*); reasons newly identified as Expectations are added to that category. Results are qualitatively similar under both: the narrower definition yields larger point estimates with wider standard errors, consistent with additional noise from the smaller reason set; the broader definition produces estimates close to the baseline. See Appendix IA.7.3.2 for prompt details.

Table IA.7: Regressions of inflation on imbalances (with expectations)

	All		*FOMC*		*Staff*	
	Core	*Headline*	*Core*	*Headline*	*Core*	*Headline*
I_d	0.093**	0.191***	0.119***	0.198***	0.035	0.081
	(0.043)	(0.071)	(0.038)	(0.059)	(0.024)	(0.053)
I_s	0.199***	0.373***	0.126***	0.221***	0.128***	0.265***
	(0.054)	(0.085)	(0.035)	(0.063)	(0.035)	(0.060)
I_e	0.186*	0.282	0.205***	0.429***	0.012	-0.063
	(0.111)	(0.174)	(0.068)	(0.127)	(0.057)	(0.095)
π_{t-2}	0.976***	0.956***	0.978***	0.959***	0.985***	0.969***
	(0.018)	(0.021)	(0.019)	(0.023)	(0.018)	(0.021)
Obs.	636	636	636	636	636	636
R^2	0.979	0.957	0.979	0.955	0.978	0.955

Notes: The dependent variables are monthly levels of headline and core inflation, expressed as year-over-year percentage changes, from the FRED. Demand, supply, and expectations imbalances are defined as in Equation (15). We compute imbalances separately for FOMC members, staff, and the pooled set of reasons. Regressions are estimated by OLS with Newey–West standard errors with 8 lags (in parentheses). *, **, *** denote significance at the 10, 5, and 1 percent levels, respectively.

Table IA.8: Augmented Taylor Rule Estimates (with expectations)

	Pre-Volcker	*Volcker–GFC*	*Post-GFC*
φ_d	-0.73	3.35***	2.00
	(1.12)	(0.93)	(2.42)
φ_s	0.65	-0.42	0.02
	(0.83)	(0.71)	(2.29)
φ_e	3.28	7.12***	-3.63
	(2.62)	(2.09)	(3.38)
φ_π	0.69**	1.73***	2.80**
	(0.32)	(0.20)	(1.36)
φ_x	1.09***	0.67***	0.53
	(0.40)	(0.17)	(0.32)
ρ	0.89***	0.62***	0.91***
	(0.04)	(0.07)	(0.04)
N	269	228	140
R^2	0.918	0.905	0.962

Notes: Each row reports structural slope coefficients from the augmented generalized Taylor rule in Equation (19), for different samples: *Pre-Volcker* = 1960Q1–1979Q2; *Volcker-GFC* = 1979Q3–2007Q4; *Post-GFC* = 2008Q1–2025Q2. The dependent variable is the shadow rate from Wu and Xia (2016) whenever that rate is available (1990–2022) and the federal funds rate otherwise. φ_π is the response to core inflation lagged one month, φ_x to the log output gap; φ_d, φ_s, and φ_e to the demand, supply, and expectations textual imbalances. ρ is the interest-rate smoothing parameter. Coefficients are estimated from Eq. (20) and transformed as in Carvalho, Nechio, and Tristão (2021). We construct Newey-West standard errors (with 8 lags) from the auxiliary OLS regressions, and then apply the delta method to generate standard errors for the structural parameters. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table IA.9: Response of risk-neutral yields and term premia controlling for macro news (with expectations)

	Risk-Neutral Yields		*Term Premia*	
	1y	5y	1y	5y
I_d	0.385***	0.364***	-0.017	0.036
	(0.114)	(0.120)	(0.037)	(0.039)
I_s	0.098*	0.056	0.057**	0.089***
	(0.055)	(0.058)	(0.024)	(0.029)
I_e	-0.054	-0.020	0.072	0.050
	(0.211)	(0.223)	(0.069)	(0.073)
Obs.	287	287	287	287
R^2	0.979	0.965	0.884	0.971

Notes: The table re-estimates regressions of the risk-neutral yields and term premia on demand, supply, and expectations imbalances, with the set of macro-news predictors of Bauer and Swanson (2023*b*). Risk-neutral yields and term premia come from the Adrian, Crump, and Moench (2013) decomposition. Additionally, each regression controls for the log output gap, core inflation, and the risk-neutral yield and term premium lagged two meetings. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table IA.10: High-Frequency Monetary Policy Shock Regressions (with expectations)

	MPS	*S&P 500*	*ED1*	*ED4*	*VIXY*
I_d^{stat}	0.001	-0.015	0.001	-0.001	-0.002
	(0.006)	(0.064)	(0.004)	(0.008)	(0.003)
I_s^{stat}	0.021*	-0.234***	0.020**	0.018	0.012***
	(0.012)	(0.081)	(0.010)	(0.013)	(0.002)
I_e^{stat}	-0.001	-0.271***	-0.002	-0.003	0.008**
	(0.017)	(0.088)	(0.012)	(0.020)	(0.003)
Obs.	184	184	184	184	95
R^2	0.189	0.073	0.129	0.181	0.169

Notes: The table reports regressions of high-frequency changes in asset prices and the monetary policy shock of Bauer and Swanson (2023*b*) on the demand, supply, and expectations imbalances extracted from FOMC statements. Asset price changes are from ten minutes before to twenty minutes after the announcement. The dependent variables include the change in the S&P 500 e-mini futures, current-quarter and three-quarter-ahead ahead Eurodollar futures (ED1 and ED4), VIXY, and the monetary policy shock from Bauer and Swanson (2023*b*). VIXY is an exchange-traded fund tracking short-term VIX futures; changes in VIXY are computed from TAQ data. The control set corresponds to the macro news predictors in Bauer and Swanson (2023*b*), described in Section 4.6. Sample runs from May, 1999, to December, 2023. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

IA.3 Entity neutering and look-ahead bias

A central appeal of using LLMs in this paper is that they allow us to translate long meeting transcripts into structured, econometrically usable variables. A corresponding concern is look-ahead bias. Modern LLMs are trained on large corpora that include text written long after many of the meetings in our sample. As a result, even when the model is shown only contemporaneous minutes, it may infer the meeting's identity from contextual cues and implicitly lean on ex-post knowledge to classify the passage.

Our baseline pipeline is designed to limit this channel. First, we instruct the model to base its judgments solely on the provided text and to justify its output with explicit textual support. Second, and more importantly, we separate extraction from classification. In the first stage the model reads the full meeting text but performs only a constrained identification task (listing reasons and quoting supporting language). In the second stage the model receives only the specific reason–context pair and classifies it using sign-restriction logic. By design, the classifier never sees the full minutes, which reduces its ability to identify the underlying meeting and import information that is not in the local passage.

This subsection implements an additional, more aggressive safeguard based on the entity-neutering approach of Engelberg et al. (2025). The goal is to make each reason–context pair effectively unidentifiable with respect to meeting date (and any other uniquely identifying attributes), while preserving the economic content needed for the second-stage classification.

The procedure starts from the first-stage outputs: for each identified inflation reason we retain (i) a concise reason label and (ii) the accompanying explanation/context excerpt used to justify that label. This is the object that enters the second-stage classifier in the baseline pipeline.

We prompt an LLM to remove or replace any information that could anchor the passage to a particular meeting date or uniquely identify individuals or episodes. Concretely, the model replaces dates, named people, named events, and other sharp identifiers with generic placeholders (e.g., *year_x*, *person_x*, *event_x*), while being explicitly instructed to preserve the economic structure required for classification (supply vs. demand vs. expectations/monetary policy, and the implied directions for inflation and real activity).

Masking alone does not address memorization: even if explicit identifiers are removed, an LLM might still recognize distinctive phrases from its training data. We therefore add a paraphrasing step that rewrites the reason label and the supporting passage so that the token sequence and sentence structure are substantially altered, while keeping the core content fixed. The full prompts used are shown in Appendix IA.7.6.

After masking and paraphrasing, we run an adversarial identification test. A fresh, "omniscient" LLM is shown the neutered text and asked to guess the most likely meeting date. We treat the text as *identifiable* if the guessed date falls within a ±14 day window around the true meeting date. If the adversary succeeds, we re-run masking and paraphrasing more aggressively and repeat the adversarial test. We allow up to nine iterations, mirroring the iterative, adversarial structure emphasized by Engelberg et al. (2025). If the adversary fails at any iteration, we stop and retain that neutered output for downstream analysis; if the adversary continues to succeed through the iteration cap, we drop the corresponding observation.

Applying this procedure to our full set of first-stage reason–context pairs yields complete neutering within the iteration cap: every reason becomes unidentifiable within nine iterations, with the vast majority neutered after a single pass. We then feed the final neutered reason–context pairs into the same second-stage classifier as in the baseline pipeline, construct meeting-level counts and normalized shares exactly as in the main text, and compare these neutered measures to their raw counterparts.

Two results emerge. First, meeting-level demand and supply counts from the neutered text remain close to the baseline: correlations between raw and neutered series (split by inflation-increasing and inflation-decreasing reasons) exceed 90%. Second, repeating the key specifications with neutered variables yields qualitatively unchanged results. These findings indicate that our main conclusions are not driven by look-ahead bias from meeting re-identification or memorized phrasing.

IA.4 Supplementary Figures

Figure IA.1: Diagram for two-stage reason classification procedure

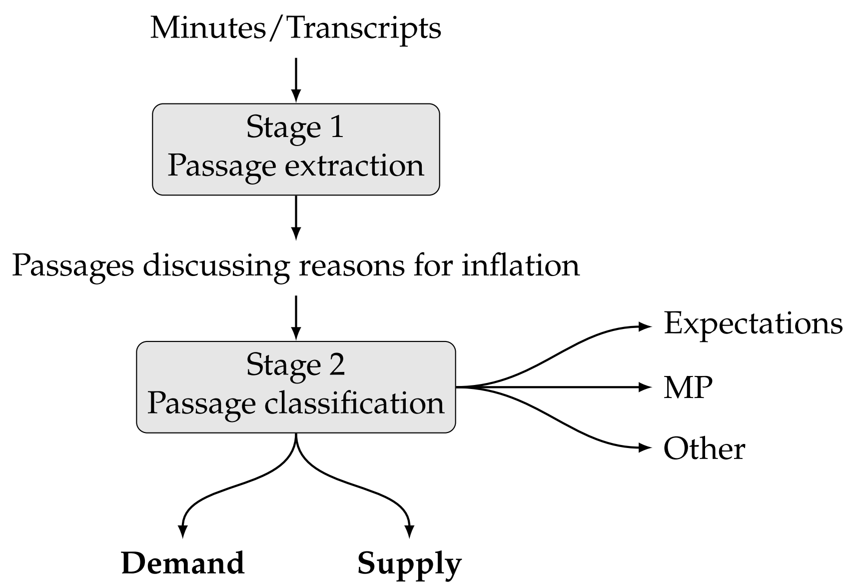

Notes: The diagram summarizes our two-stage procedure for extracting inflation reasons from FOMC textual records. Stage 1 identifies passages in the minutes where participants cite structural drivers of inflation. Stage 2 classifies reasons as demand, supply, expectations, monetary policy, or other. See the main text for more details.

Figure IA.2: Total Reason Counts Over Time

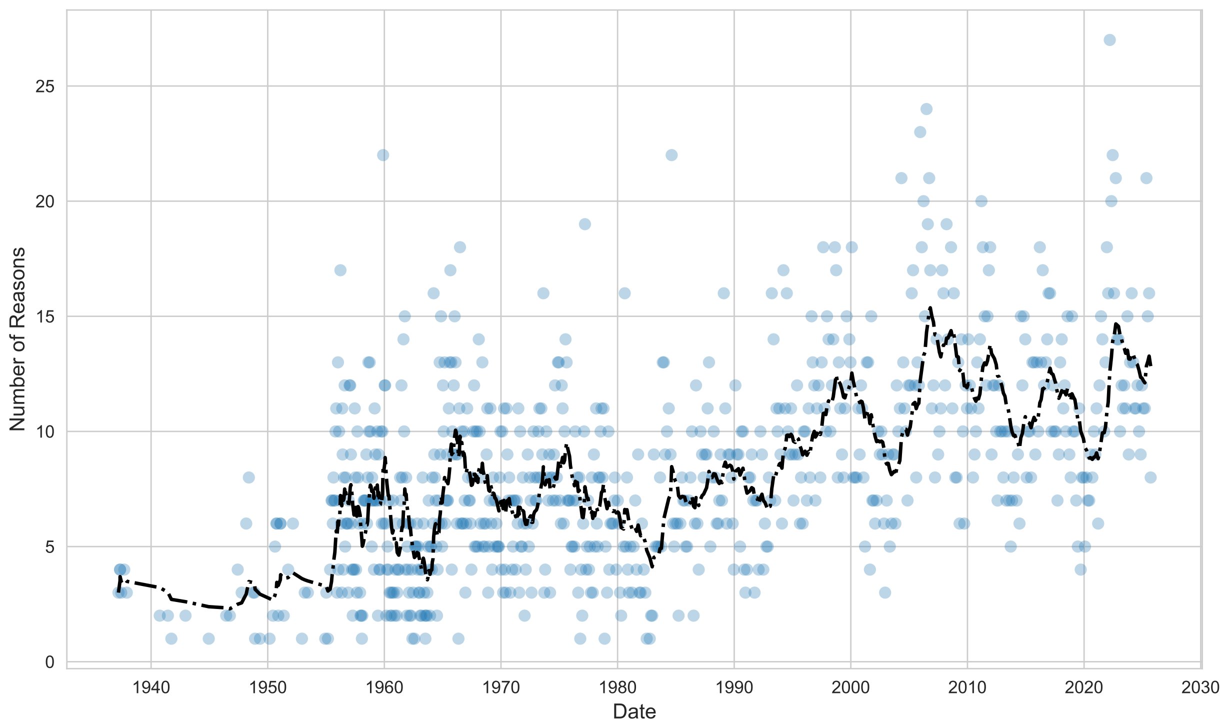

Note: Each dot gives the meeting-level count of reasons: demand, supply, or expectations, inflationary or disinflationary. The black curve is an exponential moving average with a half-life of eight meetings. See the main text for details on the construction of these variables.

Figure IA.3: Presence of demand, supply, and expectations reasons over time

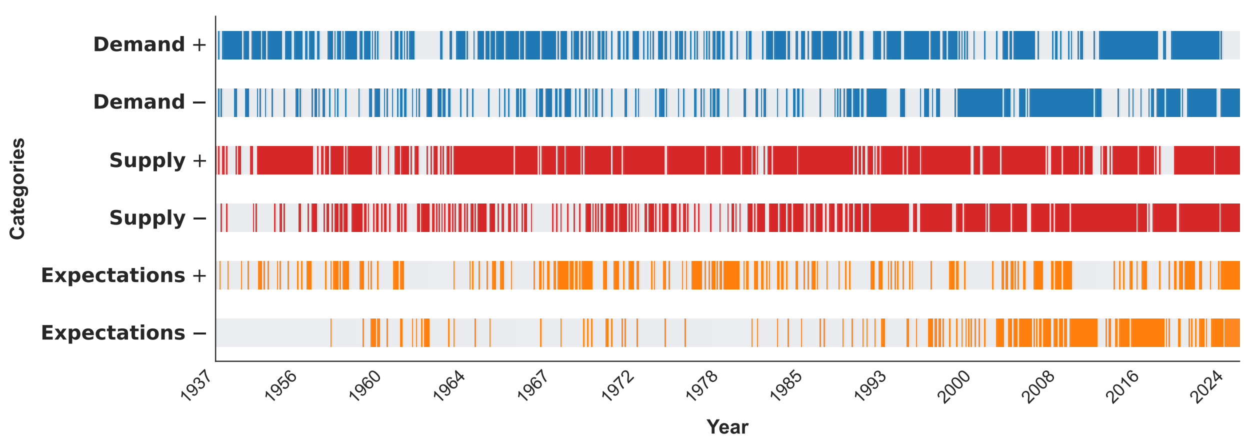

Note: A filled cell marks any FOMC meeting which mentions at least one reason of the indicated category. Inflationary (disinflationary) reasons are accompanied by a "+" ("−"). See the main text for details on the construction of the variables.

Figure IA.4: Aggregated Reason Counts Over Time (1960s)

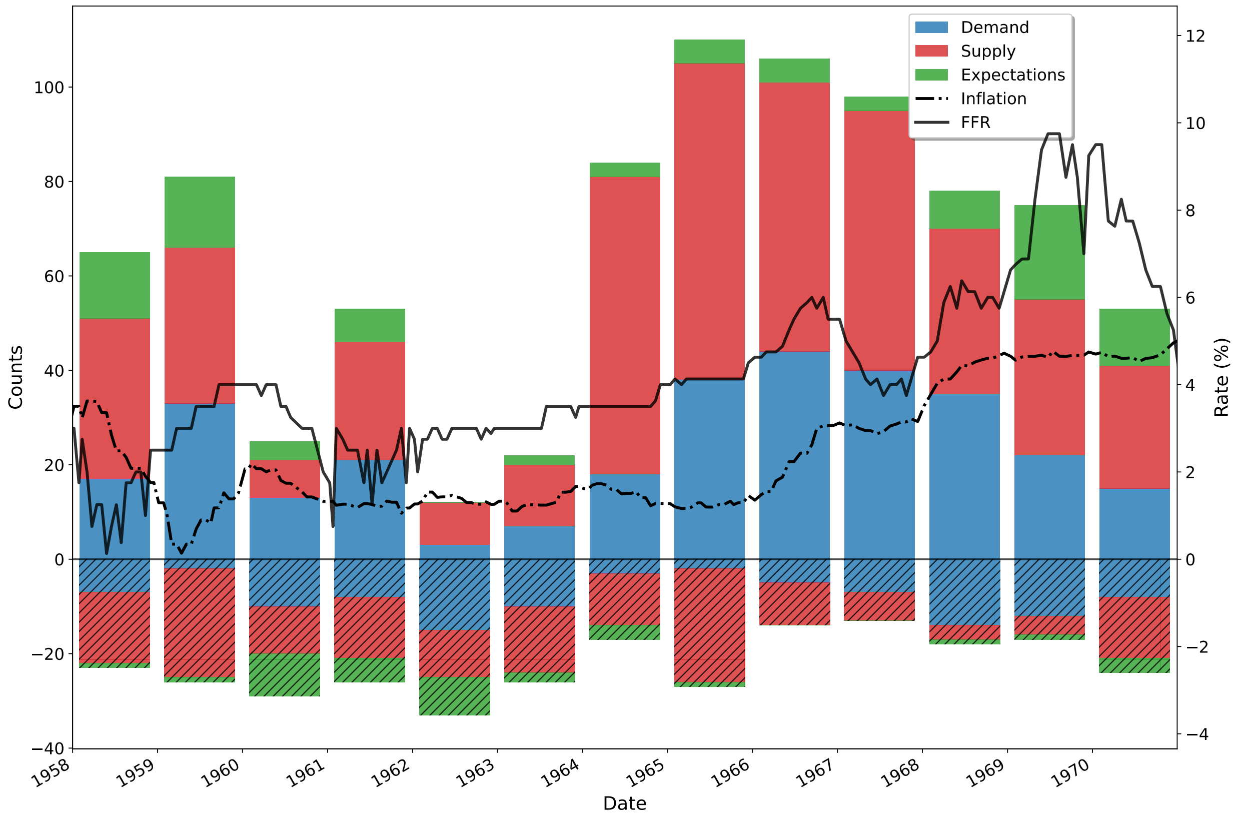

Notes: This figure aggregates annual counts of demand, supply, and expectation reasons, distinguishing inflationary (bars going up, solid) and disinflationary reasons (bars going down, hatched) on the left axis. The snapshot covers the 1960s. The Federal Funds Rate and inflation are shown on the right axis. We use CPI inflation before 1960 and core inflation afterwards. See the main text for details on variable construction.

Figure IA.5: Aggregated Reason Counts Over Time (1970s)

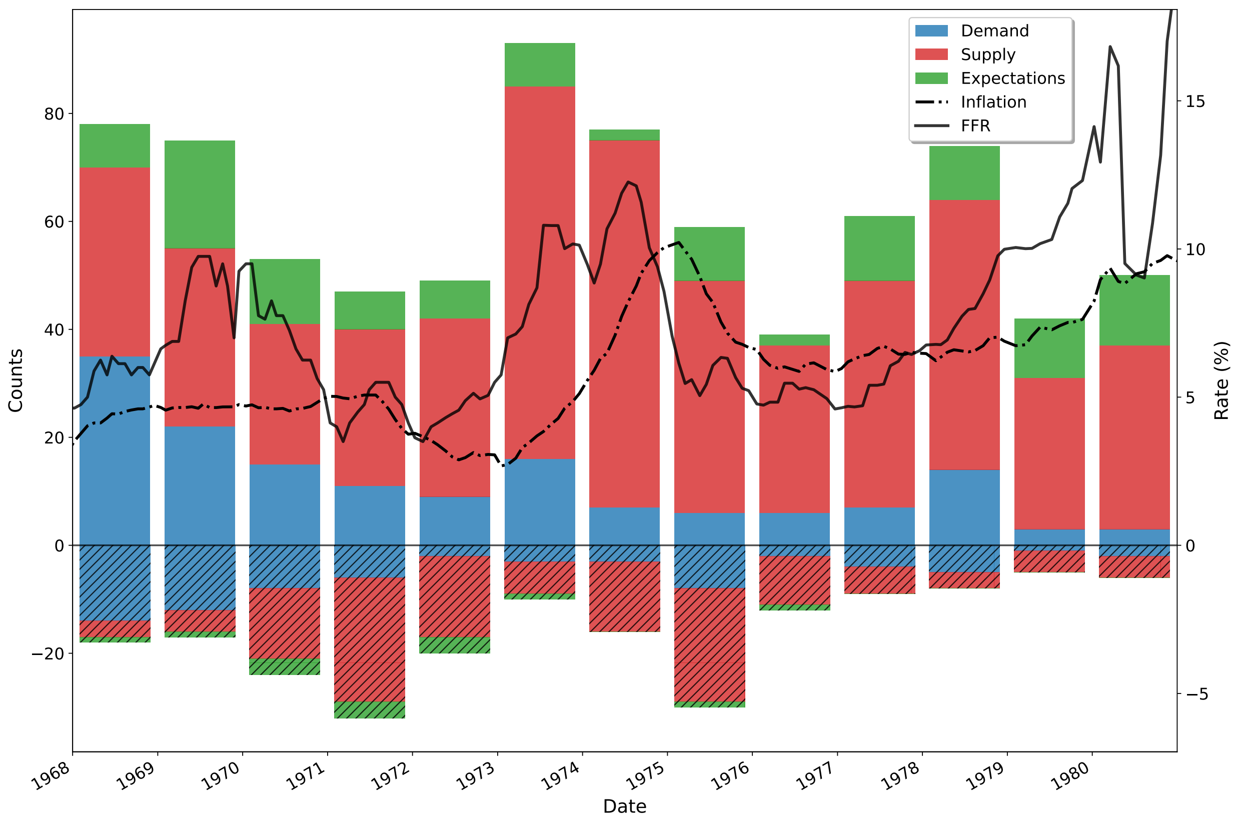

Notes: This figure aggregates annual counts of demand, supply, and expectation reasons, distinguishing inflationary (solid) and disinflationary (hatched) bars on the left axis. The snapshot covers the 1970s. The Federal Funds Rate and core inflation are shown on the right axis. See the main text for details on variable construction.

Figure IA.6: Aggregated Reason Counts Over Time (Covid)

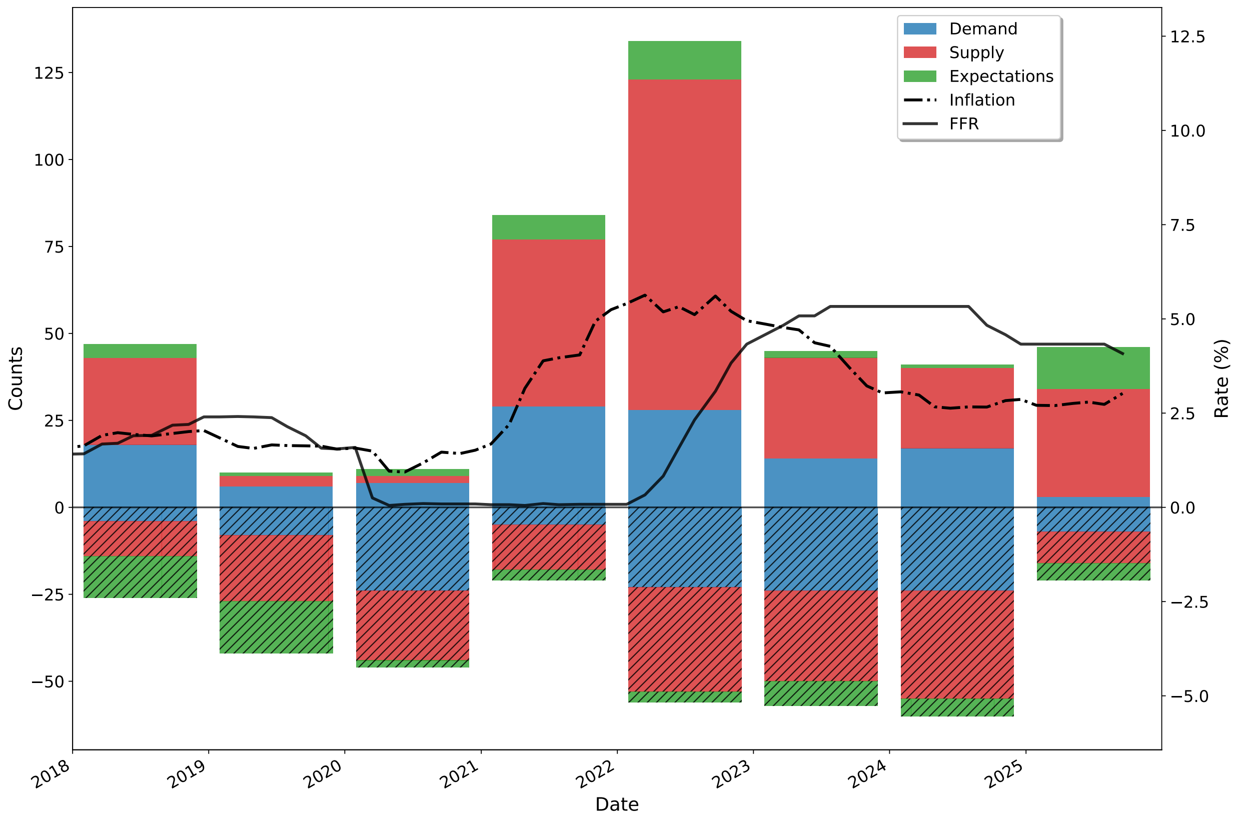

Notes: This figure aggregates annual counts of demand, supply, and expectation reasons, separating inflationary (solid) and disinflationary (hatched) bars on the left axis. The snapshot covers the Covid period. The Federal Funds Rate and core inflation are plotted on the right axis. See the main text for details on variable construction.

Figure IA.7: Share of reasons discussed by staff and FOMC members over time

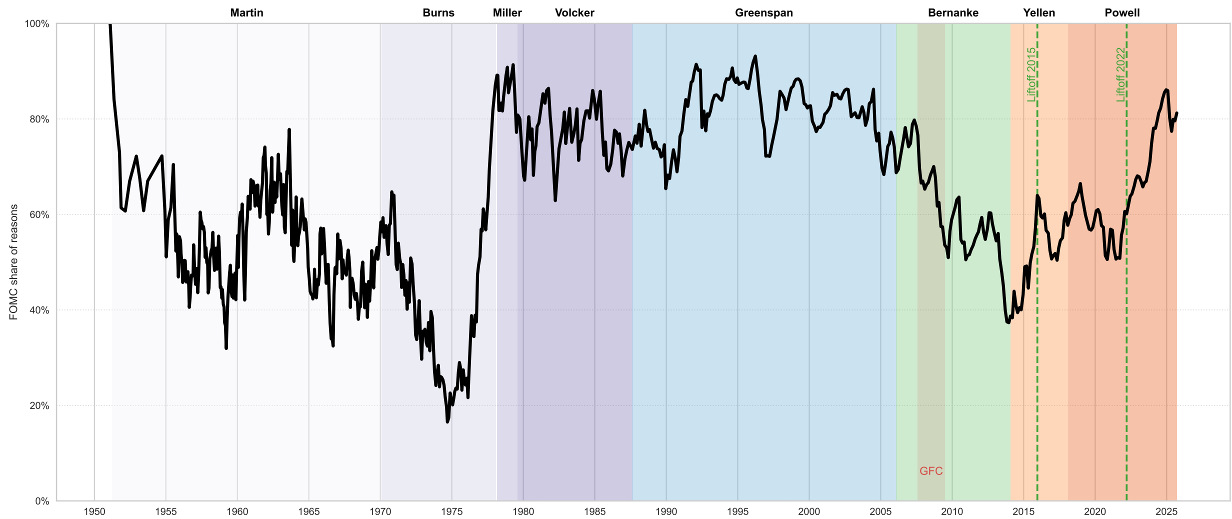

Notes: This figure plots an exponential moving average of the share of reasons presented by FOMC members in each meeting over the sample period, using a 4-meeting half-life. Shaded regions identify the tenure of each FOMC Chair.

Figure IA.8: Number of newspaper articles per year

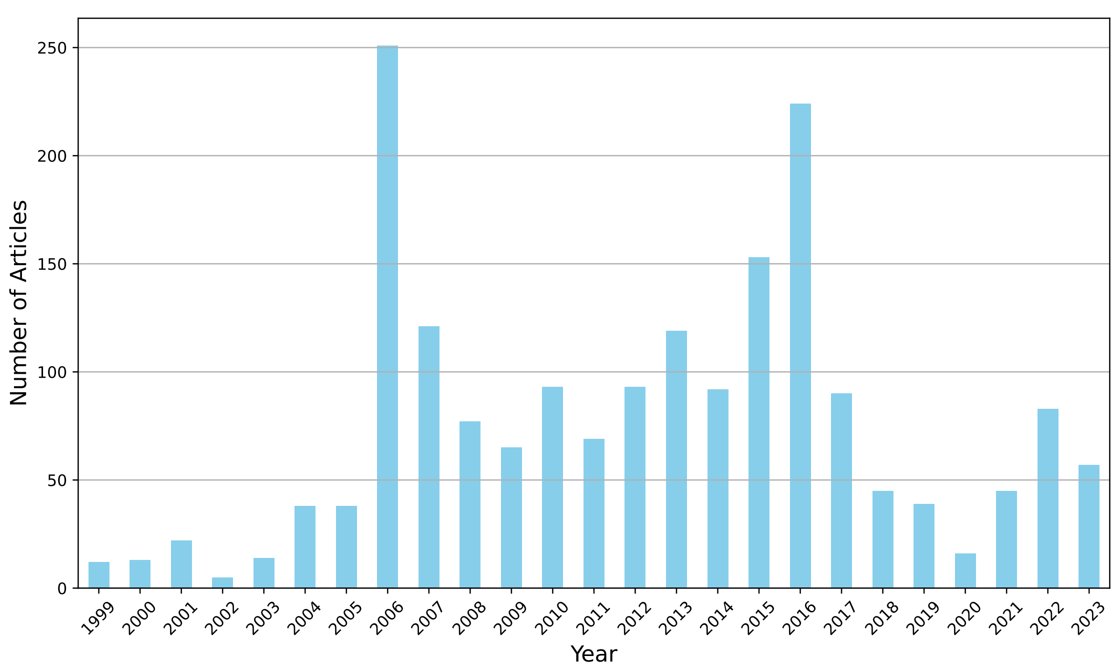

Notes: This plot presents the number of newspaper articles considered in our sample per year.

Figure IA.9: Wordcloud of the body of the newspaper articles considered

Notes: The figure presents a wordcloud from the initial filtered set of articles we consider. We lemmatize and filter tokens, removing stopwords, punctuation, and numbers. The stopword list combines the standard lists with additional domain-specific terms—both generic media boilerplate and FOMC-specific terms. Only nouns, proper nouns, and adjectives are retained.

IA.5 Supplementary Tables

Table IA.11: Movements in risk-neutral yields and term premia following the Fed's diagnosis ($t + 15$ business days)

	Risk-Neutral Yields		*Term Premia*	
	1y	*5y*	*1y*	*5y*
I_d	0.626***	0.610***	-0.108*	-0.219**
	(0.225)	(0.184)	(0.061)	(0.101)
I_s	0.066	0.090	0.103***	0.176***
	(0.110)	(0.077)	(0.039)	(0.063)
π_t	0.351***	0.249***	-0.013	-0.018
	(0.100)	(0.055)	(0.025)	(0.026)
x_t	0.174***	0.120***	-0.017*	-0.043**
	(0.047)	(0.034)	(0.010)	(0.019)
$rn_{t-2}^{(n)}$	0.750***	0.779***	0.029**	0.058***
	(0.058)	(0.041)	(0.013)	(0.021)
$tp_{t-2}^{(n)}$	0.646***	0.170***	0.723***	0.844***
	(0.219)	(0.064)	(0.070)	(0.039)
Obs.	367	367	367	367
R^2	0.943	0.949	0.779	0.896

Notes: Each column reports OLS estimates of risk-neutral yields $rn_t^{(n)}$ or term premia $tp_t^{(n)}$ on demand and supply imbalances (I_d, I_s), contemporaneous core inflation (π_t) and log output gap (x_t), and the risk-neutral yields and term premia lagged by two meetings, analogous to the auxiliary regression in Equation (20). Maturities are one and five years. Risk-neutral yields and term premia come from Adrian, Crump, and Moench (2013); yields and premia are measured fifteen business days after FOMC meeting t, including the release of the FOMC minutes. The sample spans August, 1979 to June, 2025 at the FOMC-meeting frequency. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table IA.12: Response of risk-neutral yields and term premia controlling for macro news (Post-GFC)

	Risk-Neutral Yields		*Term Premia*	
	1y	5y	1y	5y
I_d	0.150	0.206*	-0.104**	0.005
	(0.120)	(0.123)	(0.049)	(0.054)
I_s	0.048	-0.011	0.135***	0.170**
	(0.108)	(0.111)	(0.037)	(0.065)
Obs.	127	127	127	127
R^2	0.952	0.940	0.836	0.955
Controls	✓	✓	✓	✓

Notes: The table re-estimates regressions of the risk-neutral yields and term premia on demand and supply imbalances, with the set of macro news predictors of Bauer and Swanson (2023*b*). The sample is restricted to the post-GFC period (January 2008 to December 2023). Risk-neutral yields and term premia come from the Adrian, Crump, and Moench (2013) decomposition. Additionally, each regression controls for the log output gap, core inflation, and the risk-neutral yield and term premium lagged two meetings. See the main text for details on the variable definitions. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table IA.13: Regressions of statement and minutes imbalances on minutes and news imbalances

	Statement Imbalances				Minutes Imbalances	
	I_d^{stat}		I_s^{stat}		I_d	I_s
	(1)	(2)	(3)	(4)	(5)	(6)
I_d	0.877***		-0.180			
	(0.285)		(0.126)			
I_s	0.417**		0.410***			
	(0.180)		(0.102)			
I_d^{news}		0.403***		0.054	0.199***	0.120**
		(0.070)		(0.048)	(0.025)	(0.053)
I_s^{news}		-0.027		0.502***	0.001	0.434***
		(0.190)		(0.118)	(0.080)	(0.095)
Obs.	151	151	151	151	151	151
R^2	0.238	0.193	0.114	0.115	0.260	0.172

Notes: The table presents regression results of demand and supply imbalances constructed from statements (I_d^{stat} and I_s^{stat}) on the corresponding imbalances constructed from minutes and transcripts (I_d and I_s) and news articles (I_d^{news} and I_s^{news}). Columns (5)–(6) regress minutes imbalances on news imbalances. The sample is restricted to the 151 FOMC meetings (May 1999–December 2023) for which all three sources are available. Newey–West standard errors with 8 lags in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Table IA.14: Regressions of high-frequency asset price changes on the difference between statement and news imbalances

	MPS	*S&P 500*	*ED1*	*ED4*	*VIXY*
$I_d^{stat} - I_d^{news}$	-0.004	0.014	0.003	-0.010*	-0.004
	(0.005)	(0.058)	(0.005)	(0.006)	(0.003)
$I_s^{stat} - I_s^{news}$	0.023*	-0.315***	0.025*	0.017	0.011***
	(0.013)	(0.102)	(0.013)	(0.013)	(0.003)
Obs.	151	151	151	151	87
R^2	0.227	0.074	0.166	0.224	0.157
Controls	✓	✓	✓	✓	✓

Notes: The table reports regressions of high-frequency changes in asset prices and the monetary policy shock of Bauer and Swanson (2023*b*) on the difference between demand and supply imbalances extracted from FOMC statements and the corresponding imbalances extracted from news articles. Asset price changes are from ten minutes before to twenty minutes after the announcement. The dependent variables include the change in the S&P 500 e-mini futures, current and three-quarter ahead Eurodollar futures (ED1 and ED4), VIXY, and the monetary policy shock from Bauer and Swanson (2023*b*). VIXY is an exchange-traded fund tracking short-term VIX futures; changes in VIXY are computed from TAQ data. The control set corresponds to the macro news predictors in Bauer and Swanson (2023*b*), described in Section 4.6. The details about the news articles are given in Appendix IA.7.5. Sample runs from May, 1999, to June, 2023. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table IA.15: Regressions of high-frequency asset price changes on statement and news imbalances

	MPS	*S&P 500*	*ED1*	*ED4*	*VIXY*
I_d^{stat}	-0.005	0.013	-0.002	-0.009	-0.001
	(0.006)	(0.087)	(0.005)	(0.009)	(0.004)
I_s^{stat}	0.025*	-0.315***	0.026**	0.018	0.012***
	(0.014)	(0.104)	(0.013)	(0.014)	(0.003)
I_d^{news}	0.004	-0.015	-0.005	0.010	0.004
	(0.007)	(0.067)	(0.006)	(0.007)	(0.003)
I_s^{news}	-0.018	0.315**	-0.022	-0.011	-0.003
	(0.016)	(0.149)	(0.016)	(0.018)	(0.009)
Obs.	151	151	151	151	87
R^2	0.228	0.074	0.171	0.225	0.172
Controls	✓	✓	✓	✓	✓

Notes: The table reports regressions of high-frequency changes in asset prices and the monetary policy shock of Bauer and Swanson (2023*b*) on the demand and supply imbalances extracted from FOMC statements, controlling for the corresponding imbalances extracted from news articles. Asset price changes are from ten minutes before to twenty minutes after the announcement. The dependent variables include the change in the S&P 500 e-mini futures, current and three-quarter ahead Eurodollar futures (ED1 and ED4), VIXY, and the monetary policy shock from Bauer and Swanson (2023*b*). VIXY is an exchange-traded fund tracking short-term VIX futures; changes in VIXY are computed from TAQ data. The control set corresponds to the macro news predictors in Bauer and Swanson (2023*b*), described in Section 4.6. The details about the news articles are given in Appendix IA.7.5. Sample runs from May, 1999, to June, 2023. Parentheses report Newey–West standard errors (8 lags). *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

Table IA.16: Response of risk-neutral yields and term premia to imbalances, controlling for macro news and news imbalances

	Risk-Neutral Yields		*Term Premia*	
	1y	5y	1y	5y
I_d	0.224	0.151	-0.141***	-0.039
	(0.137)	(0.101)	(0.044)	(0.059)
I_s	-0.016	-0.038	0.109***	0.146***
	(0.074)	(0.084)	(0.038)	(0.054)
I_d^{news}	0.247***	0.235***	0.022	0.037
	(0.047)	(0.043)	(0.019)	(0.027)
I_s^{news}	0.080	0.134	-0.058**	-0.039
	(0.153)	(0.140)	(0.024)	(0.035)
Obs.	151	151	151	151
R^2	0.977	0.969	0.852	0.953
Controls	✓	✓	✓	✓

Notes: The table re-estimates regressions of risk-neutral yields and term premia on demand and supply imbalances, adding news-based demand and supply imbalances (I_d^{news} and I_s^{news}) as additional controls alongside the macro-news predictors of Bauer and Swanson (2023*b*). The sample is restricted to meetings with news data available (May 1999–December 2023). Risk-neutral yields and term premia come from the Adrian, Crump, and Moench (2013) decomposition. Each regression also controls for the log output gap, core inflation, $rn_{t-2}^{(n)}$, and $tp_{t-2}^{(n)}$. Newey–West standard errors with 8 lags in parentheses. *, **, and *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$.

IA.6 Output Examples

This appendix reports the full LLM classification output for four illustrative reasons discussed in the main text: from the meetings on September 16, 2008; June 15, 2022; a late-1960s meeting; and a 1970s meeting during the oil shock.

IA.6.1 Example 1: Day After Lehman's Collapse (2008-09-16)

Passage: Participants "generally thought that the outlook for inflation had improved, mainly reflecting … the weakening of the labor market."

Importance score: 6.0

Speaker: FOMC

All reasons: Y

Reason (label): weakening labor market

Category: Demand

Inflation effect: Down

Output effect: Down

Scratch: weak labor market $\Rightarrow$ cooler demand $\Rightarrow \downarrow \pi$, $\downarrow y$

IA.6.2 Example 2: Covid (2022-06-15)

Passage: Participants noted that "Most agreed that risks to inflation were skewed to the upside and cited several such risks, including those associated with ongoing supply bottlenecks and rising energy and commodity prices."

Importance score: 7.0

Speaker: FOMC

All reasons: Y

Reason (label): ongoing supply bottlenecks

Category: Supply

Inflation effect: Up

Output effect: Down

Scratch: supply bottlenecks constrain output/costs $\Rightarrow\uparrow \pi$, $\downarrow y$

IA.6.3 Example 3: Fiscal Stimulus in Late 1969 (1969-12-16)

Passage: Staff emphasized that "Prices are continuing to rise at a rapid rate, and the forces making for further inflation appear impressively strong. The fiscal outlook is in a state of disarray, with all of the questions to be resolved concerned with how much more stimulative fiscal policy will become."

Importance score: 8.0

Speaker: Staff

All reasons: Y

Reason (label): fiscal outlook turning more stimulative

Category: Demand

Inflation effect: Up

Output effect: Up

Scratch: fiscal stimulus $\Rightarrow\uparrow$ demand $\Rightarrow\uparrow \pi$, $\uparrow y$

IA.6.4 Example 4: Materials Shortages During the Oil Shock (1973-10-16)

Passage: Staff stated that "In this kind of environment, inflationary pressures will surely remain strong. Persisting shortages in the basic materials-producing industries will make possible upward adjustments in prices to reflect increases in costs and to take advantage of any easing in Phase IV price restraints."

Importance score: 8.0

Speaker: Staff

All reasons: Y

Reason (label): basic materials shortages

Category: Supply

Inflation effect: Up

Output effect: Down

Scratch: materials shortages $\Rightarrow\uparrow$ costs $\Rightarrow\uparrow \pi, \downarrow y$

IA.7 LLM Prompts

IA.7.1 First Stage Prompt

The following prompt identifies each reason discussed by the Fed in the minutes for the cause of inflation. After the prompt, we attach the text for a single FOMC minute. We repeat this process separately for each minute in the sample. We use the GPT-5 API for this task.[32] Since GPT-5 is a reasoning model, it allows no temperature control.

We use the structured output mode of the API, supplying a JSON schema that defines each key's type and admissible values. This enforces a model response that adheres to the structure that we are imposing. For instance, `reason` must be a string, `speaker` must be one of staff, FOMC, both, or `na`, and so on. We also elicit an importance score between 0 and 10, which we use to filter reasons.

The prompt is as follows:

First Stage Prompt: Identifying Reasons

```
You are an economist analyzing the following FOMC minutes. Read the text
carefully and, based solely on its content, answer using minimal language and
the abbreviations below:

* Use "Y" for yes, "N" for no, and "na" if unclear.
* For the effect, use "I" for Increasing or "D" for Decreasing (or "na" if
unclear). The effect should capture whether the factor is perceived by FOMC
participants, staff, other speakers, or the meeting narration (e.g.,
"participants noted that", staff briefings) as contributing to an increase or a
decrease in inflation (current or expected), all else constant.
* For the speaker, use "FOMC" for FOMC participants, "staff" for staff, "both"
for both FOMC participants and staff, and "na" if unclear (e.g., the narration
does not allow identification). Heuristics:
  - Text under sections titled "Staff Review of the Economic Situation/Financial
  Situation" or phrased as "the staff reported/briefed/noted" → "staff".
```

[32] More information about the model and the API are available at https://platform.openai.com/docs/models/gpt-5.

```
  - Text under "Participants" / "Committee" sections or phrased as "participants
  noted/judged/remarked" → "FOMC".
  - If the same passage explicitly attributes to both participants and staff →
  "both".
  - If attribution cannot be inferred from the immediate context → "na".

Answer the following:

Reasons for the Deviation
- List ALL reasons that FOMC participants, staff, other speakers, or the meeting
narration explicitly cite in the text as structural causes or underlying factors
driving inflation (i.e., only those the text itself points to as
explanations/threats to inflation or changes in prices, including concerns about
future inflation). Do not include self-referential answers such as "increasing
inflation" or just "inflation concerns" without an accompanying structural cause
(e.g., "oil prices").
- Include only reasons explicitly linked in the text to inflation, inflation
expectations, or prices (e.g., explicit mentions of "inflation," "prices,"
"inflation expectations," "inflation compensation," "breakevens," "CPI," "PCE,"
"deflator," or "core inflation"). Wages/labor costs and energy/commodity prices
are valid only when the passage explicitly ties them to inflation or prices.
- **Hard inclusion rule 1:** A reason may be included only if the quoted
sentence or passage that you provide contains at least one explicit reference to
inflation, prices, or inflation expectations (e.g., "inflation", "prices",
"inflation expectations", "inflation compensation", "breakevens", "CPI", "PCE",
"deflator", "core inflation"). If this explicit reference is not present in the
quoted sentence/passage itself, exclude the reason—even if macroeconomic
reasoning could connect it to inflation.
```

```
- **Hard Inclusion rule 2:** Include only reasons explicitly linked in the text
to U.S. inflation, inflation expectations, or U.S. prices (e.g., CPI, PCE…).
Assume references to ‘inflation/prices/expectations’ are U.S. unless the
sentence explicitly names a foreign country/region; if it does, exclude the
reason unless that same sentence explicitly links the foreign development to
U.S. inflation (e.g., import prices, exchange-rate pass-through, global energy
prices feeding U.S. CPI/PCE).

- When a single sentence or passage contains multiple distinct reasons (e.g.,
oil prices AND wage increases), extract and list each as a separate entry. Do
not merge multiple reasons into one entry.

- For each reason, include:
  -> "reason": A brief name (e.g., "wage increases", "energy prices"). If there
  are no reasons, answer "na".
  -> "importance_score": A score from 0 to 10. Higher scores are more critical
  and contribute to inflation more than lower scores. Higher scores also receive
  more attention in the FOMC discussion.
  -> "effect": "I" for Increasing or "D" for Decreasing (or "na" if unclear),
  capturing whether the factor is seen as contributing to an increase or
  decrease in inflation (current or expected), all else constant.
  -> "explanation": An explanation that must include the full sentence or
  passage from the minutes that supports the reason (not just a truncated
  phrase). Wrap the quoted part in single quotes (') or escape double quotes so
  that the JSON remains valid. Provide enough context so the quote can be read
  on its own, and ensure this quoted passage itself contains the explicit
  inflation/price/expectations reference required above.
  -> "speaker": Whether the subject who presented the reason was an FOMC
  participant, the Fed staff, or both. If the speaker is not explicitly
  identified or cannot be inferred from the immediate context, use "na".
- If token/space constraints prevent listing every relevant reason, include a
field "all_reasons" with value "N". Otherwise, use "Y".
```

```
Return your answer strictly as valid JSON (do NOT wrap the JSON in any markdown
or code fences) using exactly this structure:

{ "reasons": [ { "reason": "<brief name>", "importance_score": <number from 0 to
10>, "effect": "<I/D/na>", "explanation": "<short explanation with full
supporting quote>", "speaker": "<FOMC/staff/both/na>" } ], "all_reasons":
"<Y/N>" }

Do not include any markdown formatting or extra text—only output the JSON.

Example Output:

{ "reasons": [ { "reason": "wage increases", "importance_score": 6, "effect":
"I", "explanation": "FOMC participants stated that 'rising wages were putting
upward pressure on prices, raising concerns about higher inflation over the next
several quarters'.", "speaker": "FOMC" }, { "reason": "energy prices",
"importance_score": 5, "effect": "I", "explanation": "Staff noted that 'recent
increases in energy prices have begun to pass through to consumer prices, adding
to inflation pressures'.", "speaker": "staff" } ], "all_reasons": "Y" }

Now, here is the FOMC minutes excerpt to analyze:
```

IA.7.2 Second Stage Prompt

The second stage prompt categorizes each reason identified in step 1 (Appendix IA.7.1). We purposefully withhold the full document from the LLM and instead provide it only with the label (the `reason` field from step one) and the context (`explanation` field from step one). We again use the GPT-5 API for this task. Like before, we use the structured output mode of the API, supplying a JSON schema that defines each key's type and admissible values. We attach the reason and the explanation at the end of each prompt, separately for each reason-explanation pair (i.e., each observation generated from the first stage).

This task requires the LLM to use textbook economic reasoning to identify the relationship between the inflation source, inflation, and output. As discussed in the introduction and in our theoretical model, demand and supply pressures can be distinguished by their different effects on output and inflation. The fields generated in this task are:

1. `output_effect`: The expected effect of the cause on output.

2. `inflation_effect`: The expected effect of the cause on inflation.

3. `category`: A tag which can be either "demand", "supply", "expectations", or "monetary policy", based on the sign restriction implied by the inflation and output effects and by the context. In some rare cases when both demand and supply are present in the passage, the LLM classifies the reason as "Demand & Supply"; this single reason then counts for demand and supply separately. Ambiguous or tautological reasons are labelled "other".

4. `scratch`: A causal chain linking the shock to inflation and output (if such a link exists). This serves a dual purpose: it helps the LLM reason consistently and allows the econometrician to verify the validity of the reasoning, increasing transparency. This is in line with recent literature on machine learning, which has shown that intermediary reasoning significantly helps LLMs produce complex reasoning (Wei et al., 2022).

Additionally, from previous iterations of our prompt engineering, we identified that the LLM often had issues with reasons related to labor markets and trade deficits. Therefore, we added additional examples for these and provided guidelines in other edge cases as well as contrasting examples (i.e., wage increases caused by demand versus supply). The prompt classifies a reason as monetary policy when the policy stance, action, or communication is its main driver. For expectations, we restrict classification to changes relative to prior levels, excluding mere statements that "expectations remain anchored."

The prompt is as follows:

Second Stage Prompt: Categorizing Reasons

```
You are an economist classifying a reason cited in the FOMC minutes. Use ONLY
the provided text/context (no outside info).

Goal: assign each reason to exactly ONE label from:
```

```
  Demand, Supply, Supply&Demand, Expectations, MP, Other
Also report directional effects.

———————————— CORE LOGIC ————————————
Use textual cues first. Apply sign-restriction intuition only when the passage
supports it:
- Demand shock: inflation and real activity move in the SAME direction (↑π,↑y or
↓π,↓y).
- Supply shock: inflation and real activity move in OPPOSITE directions (↑π,↓y
or ↓π,↑y).

Label PRIORITY when multiple cues appear in the same passage:
  MP (if the mechanism is policy stance/action/communication) >
  Expectations (if the object is expectations/compensation/breakevens) >
  Supply&Demand (only if BOTH sides are operative for the same market/period in
  the same sentence/adjacent clause) >
  Demand >
  Supply >
  Other.

———————————— CATEGORIES (with concise examples) ————————————
Supply
  Definition: Cost-push or production-side constraints that alter marginal costs
  or available output.
  Requires: An explicit production/terms-of-trade trigger in the SAME sentence.
  Triggers include: energy/oil, commodity/input costs, import
  prices/terms-of-trade, tariffs, FX (dollar up/down) via import prices,
                  supply-chain/bottlenecks/shortages/logistics, capacity
                  constraints, productivity, regulation, lockdowns, strikes.
  Examples:
    • “Oil prices jumped” → ↑π,↓y
    • “Tariffs raised input costs; firms passed costs to prices” → ↑π,↓y
    • “Dollar depreciation raised import prices” → ↑π,↓y
```

```
    • “Supply bottlenecks eased” → ↓π,↑y
    • “COVID lockdowns in China disrupted inputs” → ↑π,↓y
    • “Strike curtailed auto production” → ↑π,↓y

Demand
  Definition: Drivers of aggregate spending/absorption: consumption, investment,
  fiscal transfers, sentiment, credit/financial conditions, income/wages,
  overheating/resource utilization.
  Labor & wages rule: Tight labor markets / wage growth ⇒ Demand by default
  (income/spending channel),
                      UNLESS the SAME sentence explicitly ties wages to cost
                      pass-through/price-setting (then ⇒ Supply).
  Examples:
    • “Consumer spending rebounded on reopening” → ↑π,↑y
    • “Tight labor market with faster nominal wage growth” → ↑π,↑y
    • “Looser financial conditions boosted spending” → ↑π,↑y
    • “Social distancing depressed services demand” → ↓π,↓y

Supply&Demand
  Use ONLY when the passage itself states BOTH excess demand AND supply
  constraints for the same market/period in the same sentence/adjacent clause.
  Example:
    • “Demand and supply imbalances pushing up prices” → both sides operative

Expectations
  Definition: Inflation expectations/anchoring, inflation compensation,
  breakevens/swaps, or firms’ pricing plans BECAUSE expected inflation changed.
  Examples:
    • “Breakevens rose” → ↑π (mechanism: expectations)
    • “Long-run expectations are well anchored” → ↓π risk

MP
```

```
  Definition: Policy stance/action/communication (rates, QE/QT, guidance,
  balance sheet) as the mechanism affecting π/y.
  Examples:
    • “Rate hikes will cool demand” → ↓π,↓y
    • “QE eased financial conditions” → ↑π,↑y

Other (Tautology/Unclear)
  Use if the text restates inflation/price levels without a driver, is pure
  accounting/base effects, or mechanism is vague/ambiguous.
  Examples:
    • “Inflation remained elevated amid broader price pressures” (no driver) →
    na/na
    • “Base effects will lift 12-month inflation temporarily” → na/Up

———————————— EDGE-CASE OVERRIDES (clause-level) ————————————
Labor & wages:
  Higher wages/tight labor markets ⇒ Demand by default; if SAME sentence links
  wages to cost pass-through/price-setting (unit labor costs, input costs,
  “passed through to prices”), ⇒ Supply.
Pass-through:
  Treat “pass-through” as Supply ONLY when explicitly passing input costs
  (energy/imports/wages as costs) into prices in the SAME sentence.
Trade balance / BoP deficits:
  Default ⇒ Demand (absorption > output). If the sentence ties higher import
  costs/shortages/terms-of-trade to prices ⇒ Supply.
Yields vs expectations:
  “Breakevens/inflation compensation/expected inflation” ⇒ Expectations.
  Pure yields/term premium/market functioning without expectations ⇒ MP if
  clearly tied to policy; otherwise Other.

———————————— OUTPUT ————————————
Return strict JSON (no markdown, no extra text):
{
```

```
  "scratch": "<brief causal chain, max 120 chars>",
  "output_effect": "<Up | Down | na>",
  "inflation_effect": "<Up | Down | na>",
  "category": "<Demand | Supply | Supply&Demand | Expectations | MP | Other>"
}

Definitions:
- scratch: terse causal chain (e.g., "oil ↑costs ⇒ ↑π,↓y"; "tightening ⇒ ↓demand
⇒ ↓π,↓y"; "breakevens ↑ ⇒ ↑π").
- output_effect: direction of real activity if implied by the passage or clear
sign-restriction; else "na".
- inflation_effect: direction of prices if implied by the passage or clear
sign-restriction; else "na".
- category: exactly one label from the set above (case-sensitive). Supply&Demand
only if BOTH sides are explicit in the same passage.

Constraints:
- Use ONLY the provided text/context.
- Prefer explicit textual cues; if mechanism is unclear, output category
"Other".
- Enumerations are strict: any value outside the allowed set is invalid.

Below is the reason to classify and the surrounding context from the FOMC
minutes.
```

IA.7.3 Alternative Expectation Prompts

This section describes the prompt and refinements used to construct the panel of reasons linking perceived inflation to expectations, discussed further in Appendix IA.2. Appendix IA.7.3.1 implements a restrictive filter that classifies a reason as expectations-based only if it concerns long-run expectations and denotes a change; mere statements that expectations are anchored do not qualify. Appendix IA.7.3.2 implements a less restrictive filter that also treats historical terminology for expectations (such as "inflation psychology") as valid reasons.

IA.7.3.1 More Restrictive Classification focusing on long-term expectation changes

Restrictive Refinement for Expectations

```
You are an economist classifying a reason cited in the FOMC minutes. Use ONLY
the provided text/context (no outside info).

Goal: assign each reason to exactly ONE label from:
  Demand, Supply, Supply&Demand, Expectations, MP, Other
Also report directional effects.

———————— CORE LOGIC ————————
Use textual cues first. Apply sign-restriction intuition only when the passage
supports it:
- Demand shock: inflation and real activity move in the SAME direction (↑π,↑y or
↓π,↓y).
- Supply shock: inflation and real activity move in OPPOSITE directions (↑π,↓y
or ↓π,↑y).

Label PRIORITY when multiple cues appear in the same passage:
  MP (if the mechanism is policy stance/action/communication) >
  Expectations (if the object is expectations/compensation/breakevens) >
  Supply&Demand (only if BOTH sides are operative for the same market/period in
  the same sentence/adjacent clause) >
  Demand >
  Supply >
  Other.

———————— CATEGORIES (with concise examples) ————————
Supply
  Definition: Cost-push or production-side constraints that alter marginal costs
  or available output.
```

```
  Requires: An explicit production/terms-of-trade trigger in the SAME sentence.
  Triggers include: energy/oil, commodity/input costs, import
  prices/terms-of-trade, tariffs, FX (dollar up/down) via import prices,
                    supply-chain/bottlenecks/shortages/logistics, capacity
                    constraints, productivity, regulation, lockdowns, strikes.
  Examples:
    • “Oil prices jumped” → ↑π,↓y
    • “Tariffs raised input costs; firms passed costs to prices” → ↑π,↓y
    • “Dollar depreciation raised import prices” → ↑π,↓y
    • “Supply bottlenecks eased” → ↓π,↑y
    • “COVID lockdowns in China disrupted inputs” → ↑π,↓y
    • “Strike curtailed auto production” → ↑π,↓y

Demand
  Definition: Drivers of aggregate spending/absorption: consumption, investment,
  fiscal transfers, sentiment, credit/financial conditions, income/wages,
  overheating/resource utilization.
  Labor & wages rule: Tight labor markets / wage growth ⇒ Demand by default
  (income/spending channel),
                      UNLESS the SAME sentence explicitly ties wages to cost
                      pass-through/price-setting (then ⇒ Supply).
  Examples:
    • “Consumer spending rebounded on reopening” → ↑π,↑y
    • “Tight labor market with faster nominal wage growth” → ↑π,↑y
    • “Looser financial conditions boosted spending” → ↑π,↑y
    • “Social distancing depressed services demand” → ↓π,↓y

Supply&Demand
  Use ONLY when the passage itself states BOTH excess demand AND supply
  constraints for the same market/period in the same sentence/adjacent clause.
  Example:
    • “Demand and supply imbalances pushing up prices” → both sides operative
```

```
Expectations (long-run only)
  Definition: Longer-run inflation expectations (horizon ≥ 5y; e.g.,
  “long-run/longer-run/long-term,”
              “5y5y forward,” “5–10y survey measure”) that CHANGED relative to
              prior readings/levels/trend.
  Scope restriction A (no change): Do NOT tag as Expectations if the passage
  merely says
              “anchored,” “stable,” “unchanged,” or “little changed” with no
              explicit movement.
  Scope restriction B (horizon filter): Do NOT tag as Expectations if the
  passage refers to short/near-term
              or year-ahead horizons (e.g., “1-year,” “12-month,” “near-term,”
              “next year,” “two-year”).
              Handle those under Demand/Supply if the mechanism is explicit;
              otherwise ⇒ Other.
  Triggers (LR-only): rose/increased/widened/drifted up;
  fell/declined/narrowed/drifted down; shifted higher/lower;
              became unanchored/re-anchored — when explicitly tied to longer-run
              measures.
  Examples:
    • “5y5y inflation compensation rose” → ↑π (mechanism: long-run expectations)
    • “Survey 5–10y inflation expectations drifted down” → ↓π
    • “Long-run expectations are well anchored” → Other (no change;
    inflation_effect = "na")
    • “One-year-ahead expectations increased” → Not Expectations (apply
    Demand/Supply/Other per mechanism)

MP
  Definition: Policy stance/action/communication (rates, QE/QT, guidance,
  balance sheet) as the mechanism affecting π/y.
  Examples:
    • “Rate hikes will cool demand” → ↓π,↓y
    • “QE eased financial conditions” → ↑π,↑y
```

```
Other (Tautology/Unclear)
  Use if the text restates inflation/price levels without a driver, is pure
  accounting/base effects, or mechanism is vague/ambiguous.
  Examples:
    • “Inflation remained elevated amid broader price pressures” (no driver) →
    na/na
    • “Base effects will lift 12-month inflation temporarily” → na/Up

──────────── EDGE-CASE OVERRIDES (clause-level) ────────────
Labor & wages:
  Higher wages/tight labor markets ⇒ Demand by default; if SAME sentence links
  wages to cost pass-through/price-setting (unit labor costs, input costs,
  “passed through to prices”), ⇒ Supply.
Pass-through:
  Treat “pass-through” as Supply ONLY when explicitly passing input costs
  (energy/imports/wages as costs) into prices in the SAME sentence.
Trade balance / BoP deficits:
  Default ⇒ Demand (absorption > output). If the sentence ties higher import
  costs/shortages/terms-of-trade to prices ⇒ Supply.
Long-run vs short-run expectations:
  Only horizons ≥ 5y (incl. 5y5y forward, “long-run/longer-run/long-term,” 5–10y
  surveys) qualify for category "Expectations".
  Year-ahead / near-term (≤ 2y) expectations — even if they move — are NOT
  "Expectations"; classify as Demand or Supply if
  the passage specifies a spending or cost-pass-through mechanism; otherwise ⇒
  Other.
Yields vs expectations:
  “Breakevens/inflation compensation/expected inflation” ⇒ Expectations.
  Pure yields/term premium/market functioning without expectations ⇒ MP if
  clearly tied to policy; otherwise Other.

──────────── OUTPUT ────────────
```

```
Return strict JSON (no markdown, no extra text):
{
  "scratch": "<brief causal chain, max 120 chars>",
  "output_effect": "<Up | Down | na>",
  "inflation_effect": "<Up | Down | na>",
  "category": "<Demand | Supply | Supply&Demand | Expectations | MP | Other>"
}

Definitions:
- scratch: terse causal chain (e.g., "oil ↑costs ⇒ ↑π,↓y"; "tightening ⇒ ↓demand
⇒ ↓π,↓y"; "breakevens ↑ ⇒ ↑π").
- output_effect: direction of real activity if implied by the passage or clear
sign-restriction; else "na".
- inflation_effect: direction of prices if implied by the passage or clear
sign-restriction; else "na".
- category: exactly one label from the set above (case-sensitive). Supply&Demand
only if BOTH sides are explicit in the same passage.

Constraints:
- Use ONLY the provided text/context.
- Prefer explicit textual cues; if mechanism is unclear, output category
"Other".
- Enumerations are strict: any value outside the allowed set is invalid.

Below is the reason to classify and the surrounding context from the FOMC
minutes.
```

IA.7.3.2 Less Restrictive Classification Including Inflation Psychology Discussion

Capturing Broader Inflation Expectations Discussion

```
You are an economist classifying a reason cited in the FOMC minutes. Use ONLY
the provided text/context (no outside info).

Goal: assign each reason to exactly ONE label from:
  Demand, Supply, Supply&Demand, Expectations, MP, Other
Also report directional effects.

———————————— CORE LOGIC ————————————
Use textual cues first. Apply sign-restriction intuition only when the passage
supports it:
- Demand shock: inflation and real activity move in the SAME direction (↑π,↑y or
↓π,↓y).
- Supply shock: inflation and real activity move in OPPOSITE directions (↑π,↓y
or ↓π,↑y).

Label PRIORITY when multiple cues appear in the same passage:
  MP (if the mechanism is policy stance/action/communication) >
  Expectations (if the object is inflation expectations/inflation
  psychology/compensation/breakevens) >
  Supply&Demand (only if BOTH sides are operative for the same market/period in
  the same sentence/adjacent clause) >
  Demand >
  Supply >
  Other.

———————————— CATEGORIES (with concise examples) ————————————
Supply
  Definition: Cost-push or production-side constraints that alter marginal costs
  or available output.
  Requires: An explicit production/terms-of-trade trigger in the SAME sentence.
```

```
  Triggers include: energy/oil, commodity/input costs, import
  prices/terms-of-trade, tariffs, FX (dollar up/down) via import prices,
                  supply-chain/bottlenecks/shortages/logistics, capacity
                  constraints, productivity, regulation, lockdowns, strikes.
  Examples:
    • "Oil prices jumped" → ↑π,↓y
    • "Tariffs raised input costs; firms passed costs to prices" → ↑π,↓y
    • "Dollar depreciation raised import prices" → ↑π,↓y
    • "Supply bottlenecks eased" → ↓π,↑y
    • "COVID lockdowns in China disrupted inputs" → ↑π,↓y
    • "Strike curtailed auto production" → ↑π,↓y

Demand
  Definition: Drivers of aggregate spending/absorption: consumption, investment,
  fiscal transfers, sentiment, credit/financial conditions, income/wages,
  overheating/resource utilization.
  Labor & wages rule: Tight labor markets / wage growth ⇒ Demand by default
  (income/spending channel),
                    UNLESS the SAME sentence explicitly ties wages to cost
                    pass-through/price-setting (then ⇒ Supply),
                    OR explicitly ties wages to inflation
                    expectations/psychology (then ⇒ Expectations).
  Examples:
    • "Consumer spending rebounded on reopening" → ↑π,↑y
    • "Tight labor market with faster nominal wage growth" → ↑π,↑y
    • "Looser financial conditions boosted spending" → ↑π,↑y
    • "Social distancing depressed services demand" → ↓π,↓y

Supply&Demand
  Use ONLY when the passage itself states BOTH excess demand AND supply
  constraints for the same market/period in the same sentence/adjacent clause.
  Example:
    • "Demand and supply imbalances pushing up prices" → both sides operative
```

```
Expectations
  Definition: Inflation expectations or “inflationary psychology”
  (beliefs/attitudes about future inflation),
              typically longer-run (e.g., “long-run/longer-run/long-term,” “5y5y
              forward,” “5–10y survey measure”)
              or NOT explicitly tied to a short horizon, that CHANGED relative
              to prior readings/levels/trend.
  Scope restriction A (anchored/no-change): If the passage merely says
              “anchored,” “stable,” “unchanged,” “little changed,” “no inflation
              psychology,” or
              “no evidence of inflationary expectations” with no explicit
              movement in inflation or activity,
              still tag as Expectations but set inflation_effect = "na" and
              output_effect = "na"
              (do NOT infer a downward inflation effect from anchoring/absence
              alone).
  Scope restriction B (horizon filter): Do NOT tag as Expectations if the
  passage explicitly refers to short/near-term
              or year-ahead horizons (e.g., “1-year,” “12-month,” “near-term,”
              “next year,” “two-year,” “next few quarters”).
              Handle those under Demand/Supply if the mechanism is explicit;
              otherwise ⇒ Other.
  Triggers: rose/increased/widened/drifted up; fell/declined/narrowed/drifted
  down; shifted higher/lower;
              became unanchored/re-anchored — when explicitly tied to inflation
              expectations or inflationary psychology.
  Examples:
    • “5y5y inflation compensation rose” → ↑π (mechanism: expectations)
    • “Survey 5–10y inflation expectations drifted down” → ↓π
    • “Long-run expectations are well anchored” → Expectations (no change;
    inflation_effect = "na")
```

```
    • “One-year-ahead expectations increased” → Not Expectations (apply
    Demand/Supply/Other per mechanism)
    • “Inflationary psychology has strengthened” → ↑π (mechanism: expectations)

MP
  Definition: Policy stance/action/communication (rates, QE/QT, guidance,
  balance sheet) as the mechanism affecting π/y.
  Examples:
    • “Rate hikes will cool demand” → ↓π,↓y
    • “QE eased financial conditions” → ↑π,↑y

Other (Tautology/Unclear)
  Use if the text restates inflation/price levels without a driver, is pure
  accounting/base effects, or mechanism is vague/ambiguous.
  Examples:
    • “Inflation remained elevated amid broader price pressures” (no driver) →
    na/na
    • “Base effects will lift 12-month inflation temporarily” → na/Up

──────────── EDGE-CASE OVERRIDES (clause-level) ────────────
Labor & wages:
  Higher wages/tight labor markets ⇒ Demand by default; if SAME sentence links
  wages to cost pass-through/price-setting (unit labor costs, input costs,
  “passed through to prices”), ⇒ Supply; if SAME sentence links wages to
  inflation expectations/psychology, ⇒ Expectations.
Pass-through:
  Treat “pass-through” as Supply ONLY when explicitly passing input costs
  (energy/imports/wages as costs) into prices in the SAME sentence.
Trade balance / BoP deficits:
  Default ⇒ Demand (absorption > output). If the sentence ties higher import
  costs/shortages/terms-of-trade to prices ⇒ Supply.
Long-run vs short-run expectations:
```

```
  Treat inflation expectations or inflationary psychology as Expectations by
  default when the horizon is not explicitly short (≤ 2y).
  Year-ahead / near-term (≤ 2y) expectations — when explicitly mentioned (e.g.,
  “1-year,” “12-month,” “next year,” “next few quarters,” “near-term,”
  “short-term”) — are NOT "Expectations"; classify as Demand or Supply if
  the passage specifies a spending or cost-pass-through mechanism; otherwise ⇒
  Other.
Yields vs expectations:
  “Breakevens/inflation compensation/expected inflation” ⇒ Expectations.
  Pure yields/term premium/market functioning without expectations ⇒ MP if
  clearly tied to policy; otherwise Other.

———————————— OUTPUT ————————————
Return strict JSON (no markdown, no extra text):
{
  "scratch": "<brief causal chain, max 120 chars>",
  "output_effect": "<Up | Down | na>",
  "inflation_effect": "<Up | Down | na>",
  "category": "<Demand | Supply | Supply&Demand | Expectations | MP | Other>"
}

Definitions:
- scratch: terse causal chain (e.g., "oil ↑costs ⇒ ↑π,↓y"; "tightening ⇒ ↓demand
⇒ ↓π,↓y"; "breakevens ↑ ⇒ ↑π").
- output_effect: direction of real activity if implied by the passage or clear
sign-restriction; else "na".
- inflation_effect: direction of prices if implied by the passage or clear
sign-restriction; else "na".
- category: exactly one label from the set above (case-sensitive). Supply&Demand
only if BOTH sides are explicit in the same passage.

Constraints:
- Use ONLY the provided text/context.
```

```
- Prefer explicit textual cues; if mechanism is unclear, output category
"Other".
- Enumerations are strict: any value outside the allowed set is invalid.

Below is the reason to classify and the surrounding context from the FOMC
minutes.
```

IA.7.4 Statement Prompts

This appendix reports the prompts used for the FOMC statements. As before, we use a two-stage procedure with prompts that closely match those for the minutes. We also flag whether each reason is explicitly linked to inflation in the text or whether the link is inferred by the LLM. Because statements are much shorter, we do not filter reasons by importance and do not restrict the sample to explicit reasons in our baseline specification.

First Stage Prompt: Statements

```
You are an economist analyzing the following FOMC statement. Your task is to
extract only structural reasons that the statement links to
inflation/prices/expectations or to the output gap/slack/aggregate demand as
discussed in the statement. Use minimal language and the abbreviations below.

* Use "Y" or "N" only for the all_reasons field.

Answer the following:

Reasons for the Deviation
- List ALL reasons that the statement explicitly cites as structural causes or
underlying factors driving inflation or the output gap (positively or
negatively) as presented in this statement. Do not include self-referential
answers such as "increasing inflation" or just "inflation concerns" without an
accompanying structural cause (e.g., "oil prices").
```

```
- Include reasons explicitly linked in the text to (i) inflation, inflation
expectations, or prices (e.g., explicit mentions of "inflation," "prices,"
"inflation expectations," "inflation compensation," "breakevens," "CPI," "PCE,"
"deflator," or "core inflation"), or (ii) the output gap/slack/aggregate demand
vs potential (e.g., "slack," "capacity utilization," "overheating," "demand
above/below potential").
- **Hard inclusion rule 1 (Inflation OR Output-Gap trigger):** Include a reason
only if the quoted sentence/passage you provide contains either (a) an explicit
reference to inflation/prices/expectations, or (b) an explicit reference to the
output gap/slack/aggregate demand relative to potential.
- **Hard inclusion rule 2 (U.S. scope):** Include only reasons explicitly linked
in the text to U.S. inflation/prices/expectations or the U.S. output
gap/slack/aggregate demand. Assume U.S. unless the sentence names a foreign
country/region; if it does, exclude unless that same sentence explicitly links
it to U.S. inflation/prices or to the U.S. output gap (e.g., import prices,
exchange-rate pass-through, demand spillovers affecting U.S. slack).

- When a single sentence/passage contains multiple distinct structural causes
(e.g., oil prices AND wage increases), list each as a separate entry. Do not
merge multiple causes.
- **No double-counting for joint inflation–gap statements:** If the same
sentence/passage links an output-gap/slack concept directly to inflation for one
driver, record one reason only—use the most specific structural driver (e.g.,
"tight labor markets") and do not add a separate "output gap" reason. If the
sentence provides two distinct drivers (e.g., "tight labor markets and higher
oil prices are boosting inflation"), list both.

For each reason, include:
  -> "reason": A brief name (e.g., "wage increases", "energy prices"). Prefer
  the most specific structural driver mentioned (e.g., "tight labor markets"
  instead of "output gap").
  -> "importance_score": A score from 0 to 10 reflecting the statement’s
  emphasis on how much this driver matters at this time (not in general).
```

```
  -> "explanation": A short explanation that includes the full sentence or
  passage supporting the reason. Wrap the quoted part in single quotes (') or
  escape double quotes so JSON remains valid. Ensure the quoted passage itself
  contains the trigger from Hard inclusion rule 1.

Token constraints:
- If space prevents listing every relevant reason, set "all_reasons" = "N". If
you listed all reasons found, set "all_reasons" = "Y".
- If the statement contains no reasons that meet Hard inclusion rules 1–2,
return an empty "reasons" array and set "all_reasons" = "Y".

Return your answer strictly as valid JSON (do NOT wrap the JSON in markdown/code
fences) using exactly this structure:

{ "reasons": [ { "reason": "<brief name>", "importance_score": <number from 0 to
10>, "explanation": "<short explanation with full supporting quote>" } ],
"all_reasons": "<Y/N>" }

Example Output:

{ "reasons": [ { "reason": "tight labor markets", "importance_score": 7,
"explanation": "The statement notes that 'labor markets remain tight, which is
keeping upward pressure on prices and inflation expectations'." }, { "reason":
"energy prices", "importance_score": 5, "explanation": "It also states that
'recent increases in energy prices are expected to keep headline inflation
elevated in coming months'." } ], "all_reasons": "Y" }

Now, here is the statement to analyze:
```

Second Stage Prompt: Statements

```
You are an economist classifying a reason cited in an FOMC statement. Use ONLY
the provided text/context (no outside info).

Goal: assign each reason to exactly ONE label from:
  Demand, Supply, Supply&Demand, Expectations, MP, Other
Also report directional effects.

———————————— EXPLICITNESS FLAG (tag only; do not filter) ————————————
Add a field "explicit" ∈ {Y, N, na}:
- Set explicit="Y" if the quoted sentence/passage in the provided context
contains an explicit reference to inflation, prices, or inflation expectations
(e.g., “inflation”, “prices”, “inflation expectations”, “inflation
compensation”, “breakevens”, “CPI”, “PCE”, “deflator”, “core inflation”). If
multiple adjacent sentences form the passage, any one of them containing such a
term is sufficient.
- Set explicit="N" if the reason is only indirectly linked to inflation via
macro reasoning (e.g., output gap, overheating, slack, policy stance) and the
passage itself lacks any explicit inflation/price/expectations term.
- Set explicit="na" only if the context is too truncated/ambiguous to decide.
This flag does NOT change the category or effects—you must still classify
normally.

———————————— CORE LOGIC ————————————
Use textual cues first. Apply sign-restriction intuition only when the passage
supports it:
- Demand shock: inflation and real activity move in the SAME direction (↑π,↑y or
↓π,↓y).
- Supply shock: inflation and real activity move in OPPOSITE directions (↑π,↓y
or ↓π,↑y).

Label PRIORITY when multiple cues appear in the same passage:
  MP > Expectations > Supply&Demand > Demand > Supply > Other.
```

```
———————————— CATEGORIES (with concise examples) ————————————
Supply
  Definition: Cost-push or production-side constraints that alter marginal costs
  or available output.
  Requires: An explicit production/terms-of-trade trigger in the SAME sentence.
  Triggers include: energy/oil, commodity/input costs, import
  prices/terms-of-trade, tariffs, FX (dollar up/down) via import prices,
                    supply-chain/bottlenecks/shortages/logistics, capacity
                    constraints, productivity, regulation, lockdowns, strikes.
  Examples:
    • “Oil prices jumped” → ↑π,↓y
    • “Tariffs raised input costs; firms passed costs to prices” → ↑π,↓y
    • “Dollar depreciation raised import prices” → ↑π,↓y
    • “Supply bottlenecks eased” → ↓π,↑y
    • “COVID lockdowns in China disrupted inputs” → ↑π,↓y
    • “Strike curtailed auto production” → ↑π,↓y

Demand
  Definition: Drivers of aggregate spending/absorption: consumption, investment,
  fiscal transfers, sentiment, credit/financial conditions, income/wages,
  overheating/resource utilization.
  Labor & wages rule: Tight labor markets / wage growth ⇒ Demand by default
  (income/spending channel),
                      UNLESS the SAME sentence explicitly ties wages to cost
                      pass-through/price-setting (then ⇒ Supply).
  Examples:
    • “Consumer spending rebounded on reopening” → ↑π,↑y
    • “Tight labor market with faster nominal wage growth” → ↑π,↑y
    • “Looser financial conditions boosted spending” → ↑π,↑y
    • “Social distancing depressed services demand” → ↓π,↓y

Supply&Demand
```

```
  Use ONLY when the passage itself states BOTH excess demand AND supply
  constraints for the same market/period in the same sentence/adjacent clause.
  Example:
    • “Demand and supply imbalances pushing up prices” → both sides operative

Expectations
  Definition: Inflation expectations/anchoring, inflation compensation,
  breakevens/swaps, or firms’ pricing plans BECAUSE expected inflation changed.
  Examples:
    • “Breakevens rose” → ↑π (mechanism: expectations)
    • “Long-run expectations are well anchored” → ↓π risk

MP
  Definition: Policy stance/action/communication (rates, QE/QT, guidance,
  balance sheet) as the mechanism affecting π/y.
  Examples:
    • “Rate hikes will cool demand” → ↓π,↓y
    • “QE eased financial conditions” → ↑π,↑y

Other (Tautology/Unclear)
  Use if the text restates inflation/price levels without a driver, is pure
  accounting/base effects, or mechanism is vague/ambiguous.
  Examples:
    • “Inflation remained elevated amid broader price pressures” (no driver) →
    na/na
    • “Base effects will lift 12-month inflation temporarily” → na/Up

———————————— EDGE-CASE OVERRIDES (clause-level) ————————————
Labor & wages:
  Higher wages/tight labor markets ⇒ Demand by default; if SAME sentence links
  wages to cost pass-through/price-setting (unit labor costs, input costs,
  “passed through to prices”), ⇒ Supply.
Pass-through:
```

```
  Treat “pass-through” as Supply ONLY when explicitly passing input costs
  (energy/imports/wages as costs) into prices in the SAME sentence.
Trade balance / BoP deficits:
  Default ⇒ Demand (absorption > output). If the sentence ties higher import
  costs/shortages/terms-of-trade to prices ⇒ Supply.
Yields vs expectations:
  “Breakevens/inflation compensation/expected inflation” ⇒ Expectations.
  Pure yields/term premium/market functioning without expectations ⇒ MP if
  clearly tied to policy; otherwise Other.

———————————— OUTPUT ————————————
Return strict JSON (no markdown, no extra text):
{
  "scratch": "<brief causal chain, max 120 chars>",
  "output_effect": "<Up | Down | na>",
  "inflation_effect": "<Up | Down | na>",
  "category": "<Demand | Supply | Supply&Demand | Expectations | MP | Other>",
  "explicit": "<Y | N | na>"
}

Definitions:
- scratch: terse causal chain (e.g., "oil ↑costs ⇒ ↑π,↓y"; "tightening ⇒ ↓demand
⇒ ↓π,↓y"; "breakevens ↑ ⇒ ↑π").
- output_effect: direction of real activity if implied by the passage or clear
sign-restriction; else "na".
- inflation_effect: direction of prices if implied by the passage or clear
sign-restriction; else "na".
- category: exactly one label from the set above (case-sensitive). Supply&Demand
only if BOTH sides are explicit in the same passage.
- explicit: indicator of whether the reason for inflation is explicitly stated
or indirectly inferred by macroeconomic logic.

Constraints:
```

```
- Use ONLY the provided text/context.
- Prefer explicit textual cues; if mechanism is unclear, category="Other".
- Enumerations are strict (case-sensitive).
- The explicit flag is a tag, not a filter; classify the reason regardless of
its value.
Below is the reason to classify and the surrounding context from the FOMC
statement.
```

IA.7.5 News Articles and Prompts

The news dataset comes from NexisUni, a popular news database.[33] The data was downloaded using the NexisUni API. The search query follows Gorodnichenko, Pham, and Talavera (2023), and includes the term "FOMC" and at least one of the following terms: "rate" , "monetary", "federal funds rate", "fed funds rate". The time span considered is between May, 1999 and June, 2023. We restrict source type to include "Newspapers", "Newswires Press Releases", and "Webnews".

We further restrict our search to a set of relevant news sources. Our news sample combines general-audience outlets with specialized financial and banking publications. General outlets include Agence France Presse (English), The Telegraph, The Guardian, The Times (London), The Associated Press and The Associated Press International, USA Today, and The New York Times International Edition, which cover the FOMC for a broad readership, emphasizing macroeconomic conditions, political context, and household implications. Specialized outlets include The Bond Buyer, American Banker, and Associated Press Financial Wire, which target market participants and financial institutions and provide more detailed discussion of interest rates, funding conditions, and balance sheet implications. This mix allows us to capture both the public-facing narrative about FOMC decisions and the more technical coverage in the financial press.

We assign each article to the next FOMC date within 30 days; articles that cannot be matched this way are not considered in our analysis. We also de-duplicate articles assigned to an FOMC meeting date. We end up with 1,874 articles. Internet Appendix Figure IA.8 shows a time-series of the number of articles per year. Internet Appendix Figure

[33] Available at nexisuni.com.

IA.9 presents a wordcloud constructed from the content of these articles.

As with the minutes, we proceed in two stages. The first stage extracts reasons the article cites as structural causes of inflation or output-gap movements relevant to the forthcoming FOMC meeting; we retain only reasons with an importance score of at least 5. The second stage classifies each reason into demand, supply, expectations, monetary policy, or other. We also flag whether each reason is explicitly linked to inflation in the text or whether the link is inferred by the LLM (e.g. inferring that a change in the output gap affects inflation). In the baseline, we retain only explicitly linked reasons. We then compute reason counts contributing to an increase or decrease in inflation for each FOMC date and construct normalized imbalances as in Section 2. The prompts are shown below.

First Stage Prompt: News Articles

```
You are an economist analyzing the following news article mentioning a
forthcoming FOMC meeting. Your task is to extract only forward-looking
structural reasons that the article links to inflation/prices/expectations or to
the output gap/slack/aggregate demand for the upcoming FOMC decision. Exclude
retrospective summaries unless the text explicitly ties them to the next
meeting. Use minimal language and the abbreviations below.

* Use "Y" or "N" only for the all_reasons field.

Answer the following:

Reasons for the Deviation
- List ALL reasons that the article explicitly cites as structural causes or
underlying factors driving inflation or the output gap (positively or
negatively) that are relevant for the forthcoming/next FOMC meeting. Do not
include self-referential answers such as "increasing inflation" or just
"inflation concerns" without an accompanying structural cause (e.g., "oil
prices").
```

```
- Include reasons explicitly linked in the text to (i) inflation, inflation
expectations, or prices (e.g., explicit mentions of "inflation," "prices,"
"inflation expectations," "inflation compensation," "breakevens," "CPI," "PCE,"
"deflator," or "core inflation"), or (ii) the output gap/slack/aggregate demand
vs potential (e.g., "slack," "capacity utilization," "overheating," "demand
above/below potential").
- **Hard inclusion rule 1 (Inflation OR Output-Gap trigger):** Include a reason
only if the quoted sentence/passage you provide contains either (a) an explicit
reference to inflation/prices/expectations, or (b) an explicit reference to the
output gap/slack/aggregate demand relative to potential.
- **Hard inclusion rule 2 (U.S. scope):** Include only reasons explicitly linked
in the text to U.S. inflation/prices/expectations or the U.S. output
gap/slack/aggregate demand. Assume U.S. unless the sentence names a foreign
country/region; if it does, exclude unless that same sentence explicitly links
it to U.S. inflation/prices or to the U.S. output gap (e.g., import prices,
exchange-rate pass-through, demand spillovers affecting U.S. slack).
- **Hard inclusion rule 3 (Forward-looking filter):** The quoted
sentence/passage must itself contain a clear forward-looking anchor about the
upcoming/next FOMC decision, such as any of the following cues:
  "ahead of", "before", "going into", "in the run-up to", "at/for the [Month]
  meeting", "next meeting", "upcoming meeting", an explicit date of the next
  meeting, or modal/expectational language tied to the meeting (e.g., "will",
  "likely to", "expected to", "set to", "poised to", "markets are pricing",
  "odds of", "forecast to"). If the structural driver and the forward-looking
  anchor appear in adjacent sentences, include both sentences in the quoted
  passage so that the passage itself contains the forward anchor and the
  trigger.
- **Retrospective exclusion:** Exclude statements that only recap prior meetings
(e.g., "at the last meeting the Fed said…", "the Committee previously noted…")
unless the same sentence explicitly states implications for the next meeting
(e.g., "…which sets the stage for/raises the odds of X at the next meeting").
```

```
- When a single sentence/passage contains multiple distinct structural causes
(e.g., oil prices AND wage increases), list each as a separate entry. Do not
merge multiple causes. Do the same when a passage cites both price/inflation
channels and output-gap/slack channels—**unless** they are part of the same
causal chain for a single driver (see "No double-counting" below).
- **No double-counting for joint inflation–gap statements:** If the same
sentence/passage links an output-gap/slack concept directly to inflation for one
driver, record one reason only—use the most specific structural driver (e.g.,
"tight labor markets") and do not add a separate "output gap" reason. If the
sentence provides two distinct drivers (e.g., "tight labor markets and higher
oil prices are boosting inflation"), list both.

For each reason, include:
  -> "reason": A brief name (e.g., "wage increases", "energy prices"). Prefer
  the most specific structural driver mentioned (e.g., "tight labor markets"
  instead of "output gap").
  -> "importance_score": A score from 0 to 10 reflecting the article’s emphasis
  on how much this driver matters for the upcoming decision (not in general).
  -> "explanation": A short explanation that includes the full sentence or
  passage supporting the reason. Wrap the quoted part in single quotes (') or
  escape double quotes so JSON remains valid. Ensure the quoted passage itself
  contains (i) the forward-looking anchor about the next meeting and (ii) the
  trigger from Hard inclusion rule 1. If needed, include two adjacent sentences
  to satisfy both.

Token constraints:
- If space prevents listing every relevant forward-looking reason, set
"all_reasons" = "N". If you listed all forward-looking reasons found, set
"all_reasons" = "Y".
- If the article contains no forward-looking reasons that meet Hard inclusion
rules 1–3, return an empty "reasons" array and set "all_reasons" = "Y".
```

```
Return your answer strictly as valid JSON (do NOT wrap the JSON in markdown/code
fences) using exactly this structure:

{ "reasons": [ { "reason": "<brief name>", "importance_score": <number from 0 to
10>, "explanation": "<short explanation with full supporting quote>" } ],
"all_reasons": "<Y/N>" }

Example Output:

{ "reasons": [ { "reason": "tight labor markets", "importance_score": 7,
"explanation": "The article says 'with labor markets still tight ahead of the
September meeting, several analysts expect upward pressure on prices to persist
into Q4,' raising the odds of a hike." }, { "reason": "energy prices",
"importance_score": 5, "explanation": "It adds that 'recent increases in energy
prices are expected to keep headline CPI elevated going into the upcoming
meeting,' which officials will weigh." } ], "all_reasons": "Y" }

Now, here is the article to analyze:
```

Second Stage Prompt: News Articles

```
You are an economist classifying a reason cited in a news article about the FOMC
(often in the run-up to a meeting). Use ONLY the provided text/context (no
outside info).

Goal: assign each reason to exactly ONE label from:
  Demand, Supply, Supply&Demand, Expectations, MP, Other
Also report directional effects.

———————————— HARD INCLUSION (news) ————————————
```

```
- **Hard inclusion rule 1 (Explicit inflation reference required):** A reason
may be included only if the quoted sentence or passage that you provide contains
at least one explicit reference to inflation, prices, or inflation expectations
(e.g., "inflation", "prices", "inflation expectations", "inflation
compensation", "breakevens", "CPI", "PCE", "deflator", "core inflation"). If
this explicit reference is not present in the quoted sentence/passage itself,
classify as "Other" with output_effect="na" and inflation_effect="na". Do NOT
rely on macro inferences (e.g., Phillips curve) to bridge the link.

──────────── CORE LOGIC ────────────
Use textual cues first. Apply sign-restriction intuition only when the passage
supports it:
- Demand shock: inflation and real activity move in the SAME direction (↑π,↑y or
↓π,↓y).
- Supply shock: inflation and real activity move in OPPOSITE directions (↑π,↓y
or ↓π,↑y).

Label PRIORITY when multiple cues appear in the same passage:
  MP (if the mechanism is policy stance/action/communication) >
  Expectations (if the object is expectations/compensation/breakevens) >
  Supply&Demand (only if BOTH sides are operative for the same market/period in
  the same sentence/adjacent clause) >
  Demand >
  Supply >
  Other.

──────────── CATEGORIES (with concise examples) ────────────
Supply
  Definition: Cost-push or production-side constraints that alter marginal costs
  or available output.
  Requires: An explicit production/terms-of-trade trigger in the SAME sentence.
  Triggers include: energy/oil, commodity/input costs, import
  prices/terms-of-trade, tariffs, FX (dollar up/down) via import prices,
```

```
                    supply-chain/bottlenecks/shortages/logistics, capacity
                    constraints, productivity, regulation, lockdowns, strikes.
  Examples:
    • “Oil prices jumped” → ↑π,↓y
    • “Tariffs raised input costs; firms passed costs to prices” → ↑π,↓y
    • “Dollar depreciation raised import prices” → ↑π,↓y
    • “Supply bottlenecks eased” → ↓π,↑y
    • “COVID lockdowns in China disrupted inputs” → ↑π,↓y
    • “Strike curtailed auto production” → ↑π,↓y

Demand
  Definition: Drivers of aggregate spending/absorption: consumption, investment,
  fiscal transfers, sentiment, credit/financial conditions, income/wages,
  overheating/resource utilization.
  Labor & wages rule: Tight labor markets / wage growth ⇒ Demand by default
  (income/spending channel),
                    UNLESS the SAME sentence explicitly ties wages to cost
                    pass-through/price-setting (then ⇒ Supply).
  Examples:
    • “Consumer spending rebounded on reopening” → ↑π,↑y
    • “Tight labor market with faster nominal wage growth” → ↑π,↑y
    • “Looser financial conditions boosted spending” → ↑π,↑y
    • “Social distancing depressed services demand” → ↓π,↓y

Supply&Demand
  Use ONLY when the passage itself states BOTH excess demand AND supply
  constraints for the same market/period in the same sentence/adjacent clause.
  Example:
    • “Demand and supply imbalances pushing up prices” → both sides operative

Expectations
  Definition: Inflation expectations/anchoring, inflation compensation,
  breakevens/swaps, or firms’ pricing plans BECAUSE expected inflation changed.
```

```
  Examples:
    • “Breakevens rose” → ↑π (mechanism: expectations)
    • “Long-run expectations are well anchored” → ↓π risk

MP
  Definition: Policy stance/action/communication (rates, QE/QT, guidance,
  balance sheet) as the mechanism affecting π/y.
  Examples:
    • “Rate hikes will cool demand” → ↓π,↓y
    • “QE eased financial conditions” → ↑π,↑y

Other (Tautology/Unclear)
  Use if the text restates inflation/price levels without a driver, is pure
  accounting/base effects, or mechanism is vague/ambiguous.
  Examples:
    • “Inflation remained elevated amid broader price pressures” (no driver) →
    na/na
    • “Base effects will lift 12-month inflation temporarily” → na/Up

──────────── EDGE-CASE OVERRIDES (clause-level) ────────────
Labor & wages:
  Higher wages/tight labor markets ⇒ Demand by default; if SAME sentence links
  wages to cost pass-through/price-setting (unit labor costs, input costs,
  “passed through to prices”), ⇒ Supply.
Pass-through:
  Treat “pass-through” as Supply ONLY when explicitly passing input costs
  (energy/imports/wages as costs) into prices in the SAME sentence.
Trade balance / BoP deficits:
  Default ⇒ Demand (absorption > output). If the sentence ties higher import
  costs/shortages/terms-of-trade to prices ⇒ Supply.
Yields vs expectations:
  “Breakevens/inflation compensation/expected inflation” ⇒ Expectations.
```

```
  Pure yields/term premium/market functioning without expectations ⇒ MP if
  clearly tied to policy; otherwise Other.

———————————— OUTPUT ————————————
Return strict JSON (no markdown, no extra text):
{
  "scratch": "<brief causal chain, max 120 chars>",
  "output_effect": "<Up | Down | na>",
  "inflation_effect": "<Up | Down | na>",
  "category": "<Demand | Supply | Supply&Demand | Expectations | MP | Other>"
}

Definitions:
- scratch: terse causal chain (e.g., "oil ↑costs ⇒ ↑π,↓y"; "tightening ⇒ ↓demand
⇒ ↓π,↓y"; "breakevens ↑ ⇒ ↑π").
- output_effect: direction of real activity if implied by the passage or clear
sign-restriction; else "na".
- inflation_effect: direction of prices if implied by the passage or clear
sign-restriction; else "na".
- category: exactly one label from the set above (case-sensitive). Supply&Demand
only if BOTH sides are explicit in the same passage.

Constraints:
- Use ONLY the provided text/context.
- Prefer explicit textual cues; if mechanism is unclear, output category
"Other".
- Enumerations are strict: any value outside the allowed set is invalid.
- Apply Hard inclusion rule 1 strictly; if not satisfied, set category="Other"
and both effects="na".

Below is the reason to classify and the surrounding context from the news
article.
```

IA.7.6 Entity Neutering

This section presents the prompts used in the entity neutering exercise adapted from Engelberg et al. (2025) and discussed in Internet Appendix IA.3. There are three sets of prompts: masking, paraphrasing and guessing. Masking and paraphrasing have separate prompts for retries. The prompts below appear in the following order: first-try masking, retry masking, first-try paraphrasing, retry paraphrasing, and guessing.

Masking Prompt: First Try

```
Your role is to ANONYMIZE all text that is provided by the user.

You will receive JSON with two fields: reason and explanation.
You must anonymize both fields.

You must anonymize ANY information that could be used to identify (i) the
calendar time period
of the underlying meeting (year, month, quarter, decade, or any famous
historical episode),
or (ii) the identity of any individuals referenced.

You should also ANONYMIZE any other information which one could use to identify
the source
or to make an educated guess about when it occurred.

You MUST anonymize all identifiers, including names of individuals,
roles/titles, locations,
organizations, committee/body names, agencies, programs, policy tools, event
names, and ALL times,
years, dates and ALL numbers and percentages in the text, including units.

These should be replaced with person_x, role_x, location_x, org_x, committee_x,
agency_x,
```

```
program_x, policy_tool_x, event_x, time_x, year_x, date_x and number a, b, c,
respectively.
Also replace any website or internet links with link_x.

You should never just delete an identifier; instead, always replace it with an
anonymized analog.

CRITICAL: Preserve the economic meaning and structure needed for classification.

You MUST KEEP:
- whether the mechanism is supply vs demand vs expectations/persistence vs
monetary policy
- whether it refers to inflation vs real activity and the direction of effects
(up/down/ambiguous)
- generic horizons like "near-term", "medium-term", "longer-run", "past year",
"recent months"
- generic sources like "consumer surveys", "business surveys", "market-based
measures"
- generic actors like "consumers", "households", "firms", "participants",
"staff"

You MUST REMOVE/REPLACE ONLY identifiers:
- exact calendar dates/years/months/quarters
- names of people and named roles tied to specific individuals
- named events/episodes that pin down the time period (e.g., pandemics, wars,
crises) by name
- uniquely identifying survey/institution names (you may keep "consumer survey"
generically)
- precise numeric values (levels, percentages, amounts) unless strictly
necessary for the mechanism
```

```
Goal: someone should still be able to classify the passage correctly, but should
not be able to pin down the exact meeting date or specific named
individuals/events.

After you read and ANONYMIZE the text, you should output the anonymized text and
nothing else.

Return STRICT JSON ONLY with exactly these keys:
- reason
- explanation
No extra keys.
```

Masking Prompt: Retry

```
Your role is to ANONYMIZE all text that is provided by the user. You have
already tried to anonymize the following text and failed, leaving details which
were used to identify the text by an expert. Be more aggressive in your efforts
to anonymize and pay more attention, so that you do not leave even the most
minute details in the anonymized text.

You will receive JSON with two fields: reason and explanation.
You must anonymize both fields.

You must anonymize ANY information that could be used to identify (i) the
calendar time period
of the underlying meeting (year, month, quarter, decade, or any famous
historical episode),
or (ii) the identity of any individuals referenced.

You should also ANONYMIZE any other information which one could use to identify
the source
or to make an educated guess about when it occurred.
```

```
You MUST anonymize all identifiers, including names of individuals,
roles/titles, locations,
organizations, committee/body names, agencies, programs, policy tools, event
names, and ALL times,
years, dates and ALL numbers and percentages in the text, including units.

These should be replaced with person_x, role_x, location_x, org_x, committee_x,
agency_x,
program_x, policy_tool_x, event_x, time_x, year_x, date_x and number a, b, c,
respectively.
Also replace any website or internet links with link_x.

You should never just delete an identifier; instead, always replace it with an
anonymized analog.

CRITICAL: Preserve the economic meaning and structure needed for classification.

You MUST KEEP:
- whether the mechanism is supply vs demand vs expectations/persistence vs
monetary policy
- whether it refers to inflation vs real activity and the direction of effects
(up/down/ambiguous)
- generic horizons like "near-term", "medium-term", "longer-run", "past year",
"recent months"
- generic sources like "consumer surveys", "business surveys", "market-based
measures"
- generic actors like "consumers", "households", "firms", "participants",
"staff"

You MUST REMOVE/REPLACE ONLY identifiers:
- exact calendar dates/years/months/quarters
- names of people and named roles tied to specific individuals
```

```
- named events/episodes that pin down the time period (e.g., pandemics, wars,
crises) by name
- uniquely identifying survey/institution names (you may keep "consumer survey"
generically)
- precise numeric values (levels, percentages, amounts) unless strictly
necessary for the mechanism

Goal: someone should still be able to classify the passage correctly, but should
not be able to pin down the exact meeting date or specific named
individuals/events.

You have already failed me before, take no chances this time and be aggressive
in censoring the text. After you read and ANONYMIZE the text, you should output
the anonymized text and nothing else.

Return STRICT JSON ONLY with exactly these keys:
- reason
- explanation
No extra keys.
```

Paraphrasing Prompt: First Try

```
You are to paraphrase text ensuring that its core information content is
unchanged, such that
someone who has memorized the text could not recognize it. Make sure to replace
all unique or
identifying words and phrases. Also change the sentence structure so that the
resulting text is
completely different from the original. This includes changing quotes and other
highly recognizable
structures in the text.

You will receive JSON with two fields: reason and explanation.
```

```
You must paraphrase+anonymize both fields.

Any information which could ever be used to identify when the meeting occurred
or which individuals
are referenced MUST be anonymized. Absolutely no mention of key words related to
famous historical
events, named regimes, or other specific time anchors can be left in the text.

If you observe any names, proper nouns, aliases, acronyms, dates, agencies,
programs, policy tools,
or recognizable quotes, replace them with anonymous strings.

Anonymity is the prime objective, leave zero hints for someone who has memorized
the text!
Ensure that the paraphrased text has approximately the same number of words as
the original text.

CRITICAL: Preserve the economic meaning and structure needed for classification.

You MUST KEEP:
- whether the mechanism is supply vs demand vs expectations/persistence vs
monetary policy
- whether it refers to inflation vs real activity and the direction of effects
(up/down/ambiguous)
- generic horizons like "near-term", "medium-term", "longer-run", "past year",
"recent months"
- generic sources like "consumer surveys", "business surveys", "market-based
measures"
- generic actors like "consumers", "households", "firms", "participants",
"staff"

You MUST REMOVE/REPLACE ONLY identifiers:
- exact calendar dates/years/months/quarters
```

```
- names of people and named roles tied to specific individuals
- named events/episodes that pin down the time period (e.g., pandemics, wars,
crises) by name
- uniquely identifying survey/institution names (you may keep "consumer survey"
generically)
- precise numeric values (levels, percentages, amounts) unless strictly
necessary for the mechanism

Goal: someone should still be able to classify the passage correctly, but should
not be able to pin down the exact meeting date or specific named
individuals/events.

After you read and ANONYMIZE the text, you should output the anonymized text and
nothing else.

Return STRICT JSON ONLY with exactly these keys:
- reason
- explanation
No extra keys.
```

Paraphrasing Prompt: Retry

```
You are to paraphrase text ensuring that its core information content is
unchanged, such that
someone who has memorized the text could not recognize it. You have already
failed to paraphrase
the text to make it anonymous. Be more aggressive in your paraphrasing and
anonymizing of the text.
Make sure to replace all unique or identifying words and phrases. Also change
the sentence structure
so that the resulting text is completely different from the original. This
includes changing quotes
and other highly recognizable structures in the text.
```

```
You will receive JSON with two fields: reason and explanation.
You must paraphrase+anonymize both fields.

Any information which could ever be used to identify when the meeting occurred
or which individuals
are referenced MUST be anonymized. Absolutely no mention of key words related to
famous historical
events, named regimes, or other specific time anchors can be left in the text.

If you observe any names, proper nouns, aliases, acronyms, dates, agencies,
programs, policy tools,
or recognizable quotes, replace them with anonymous strings.

Anonymity is the prime objective, leave zero hints for someone who has memorized
the text!
Ensure that the paraphrased text has approximately the same number of words as
the original text.

CRITICAL: Preserve the economic meaning and structure needed for classification.

You MUST KEEP:
- whether the mechanism is supply vs demand vs expectations/persistence vs
monetary policy
- whether it refers to inflation vs real activity and the direction of effects
(up/down/ambiguous)
- generic horizons like "near-term", "medium-term", "longer-run", "past year",
"recent months"
- generic sources like "consumer surveys", "business surveys", "market-based
measures"
- generic actors like "consumers", "households", "firms", "participants",
"staff"
```

```
You MUST REMOVE/REPLACE ONLY identifiers:
- exact calendar dates/years/months/quarters
- names of people and named roles tied to specific individuals
- named events/episodes that pin down the time period (e.g., pandemics, wars,
crises) by name
- uniquely identifying survey/institution names (you may keep "consumer survey"
generically)
- precise numeric values (levels, percentages, amounts) unless strictly
necessary for the mechanism

Goal: someone should still be able to classify the passage correctly, but should
not be able to pin down the exact meeting date or specific named
individuals/events.

You have already failed once before, take no chances this time, be aggressive in
your effort to anonymize the text. After you read and ANONYMIZE the text, you
should output the anonymized text and nothing else.

Return STRICT JSON ONLY with exactly these keys:
- reason
- explanation
No extra keys.
```

Guessing Prompt

```
You will receive a body of text which has been anonymized. You are omniscient.
Use all of your knowledge and any clues in the text to infer the most likely
exact meeting date.

IMPORTANT:
- You MUST output a date in YYYY-MM-DD format even if you are unsure.
- If the text provides insufficient evidence, make your best guess and set a low
confidence.
```

```
- Do NOT refuse and do NOT output null.

Return STRICT JSON ONLY that matches the schema. Do not include any extra keys.
```